Weaponizing BGP Using Communities

Florian Streibelt, Franziska Lichtblau,
Robert Beverly, Cristel Pelsser, Georgios
Smaragdakis, Randy Bush, Anja Feldmann

BGP, the
Border Gateway Protocol

Core Protocol to Propagate

<u>Reachability of IP Prefixes</u>

Designed on Serviettes in 1994

Yes, Really

An IP Prefix

147.28.0.0/16

Autonomous System (AS)

An ISP or End Site

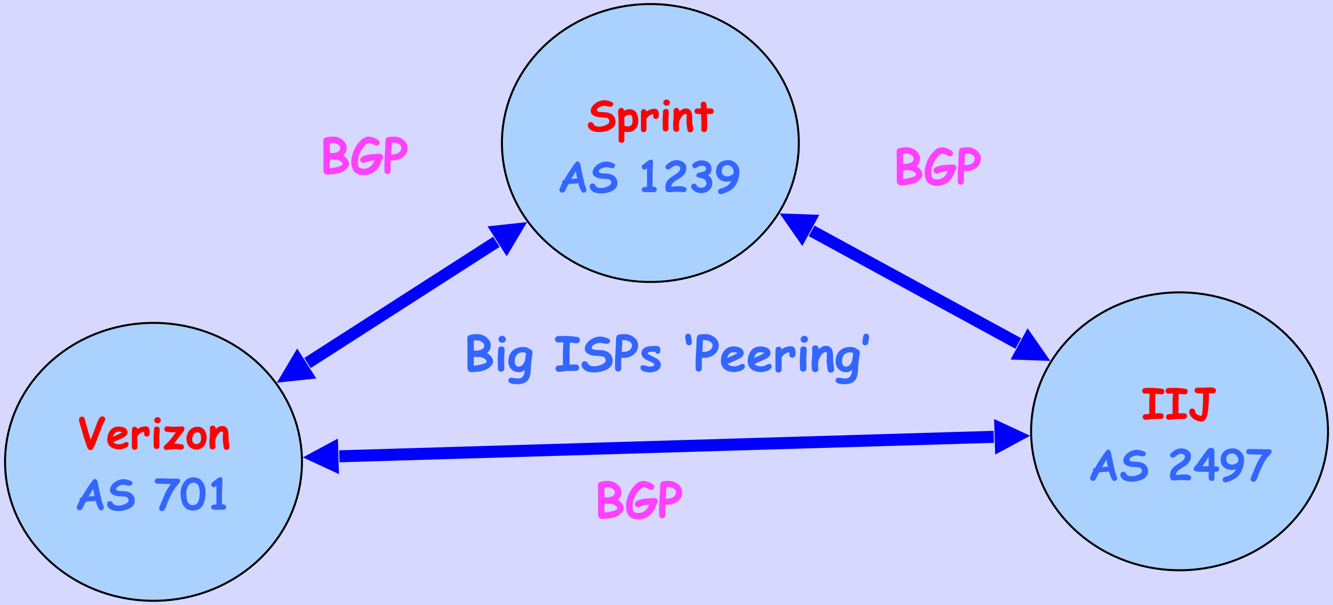

Customers

An ISP or End Site

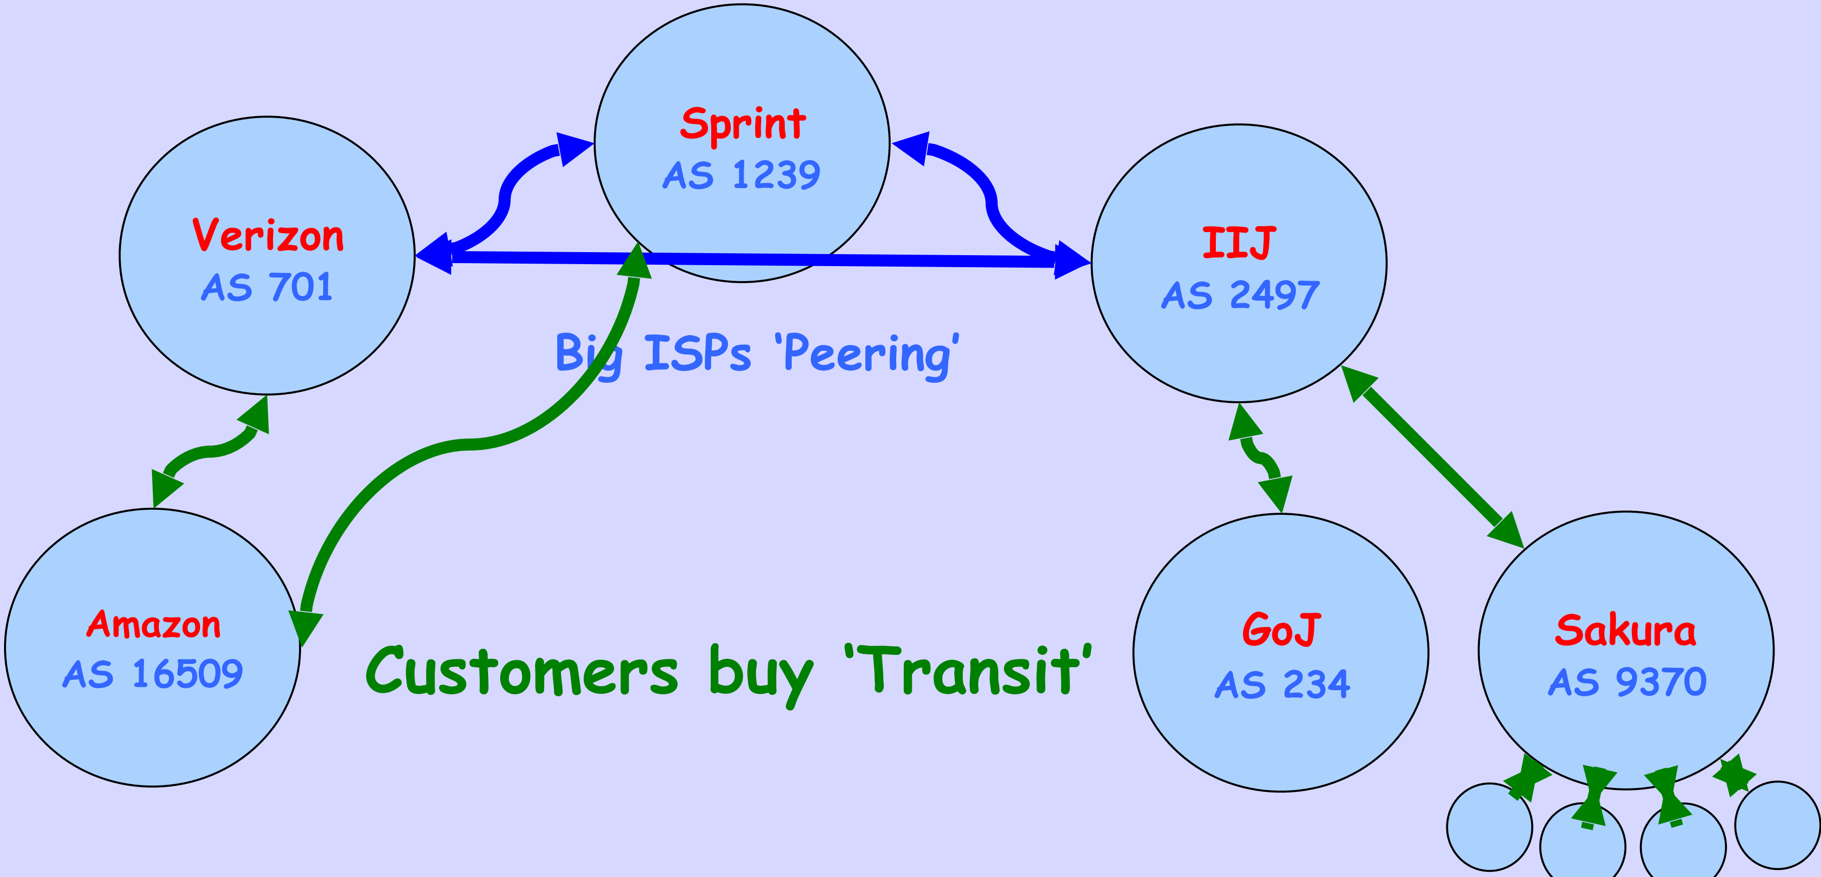

BGP Reachability Announcement

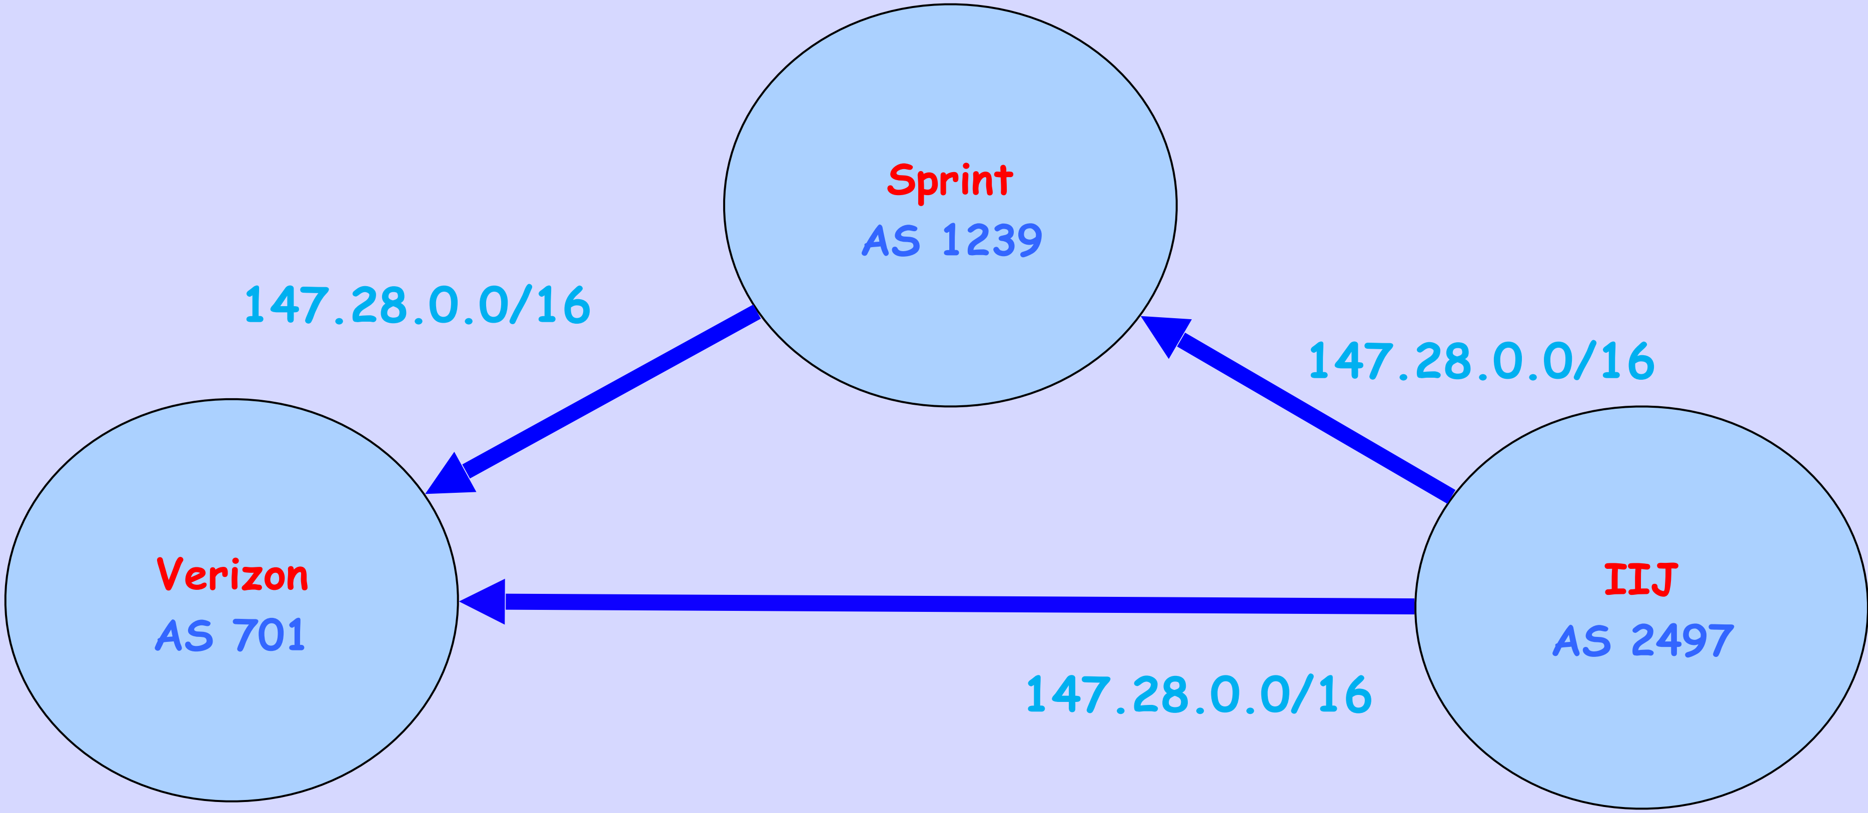

Traffic Flows Toward Announcement

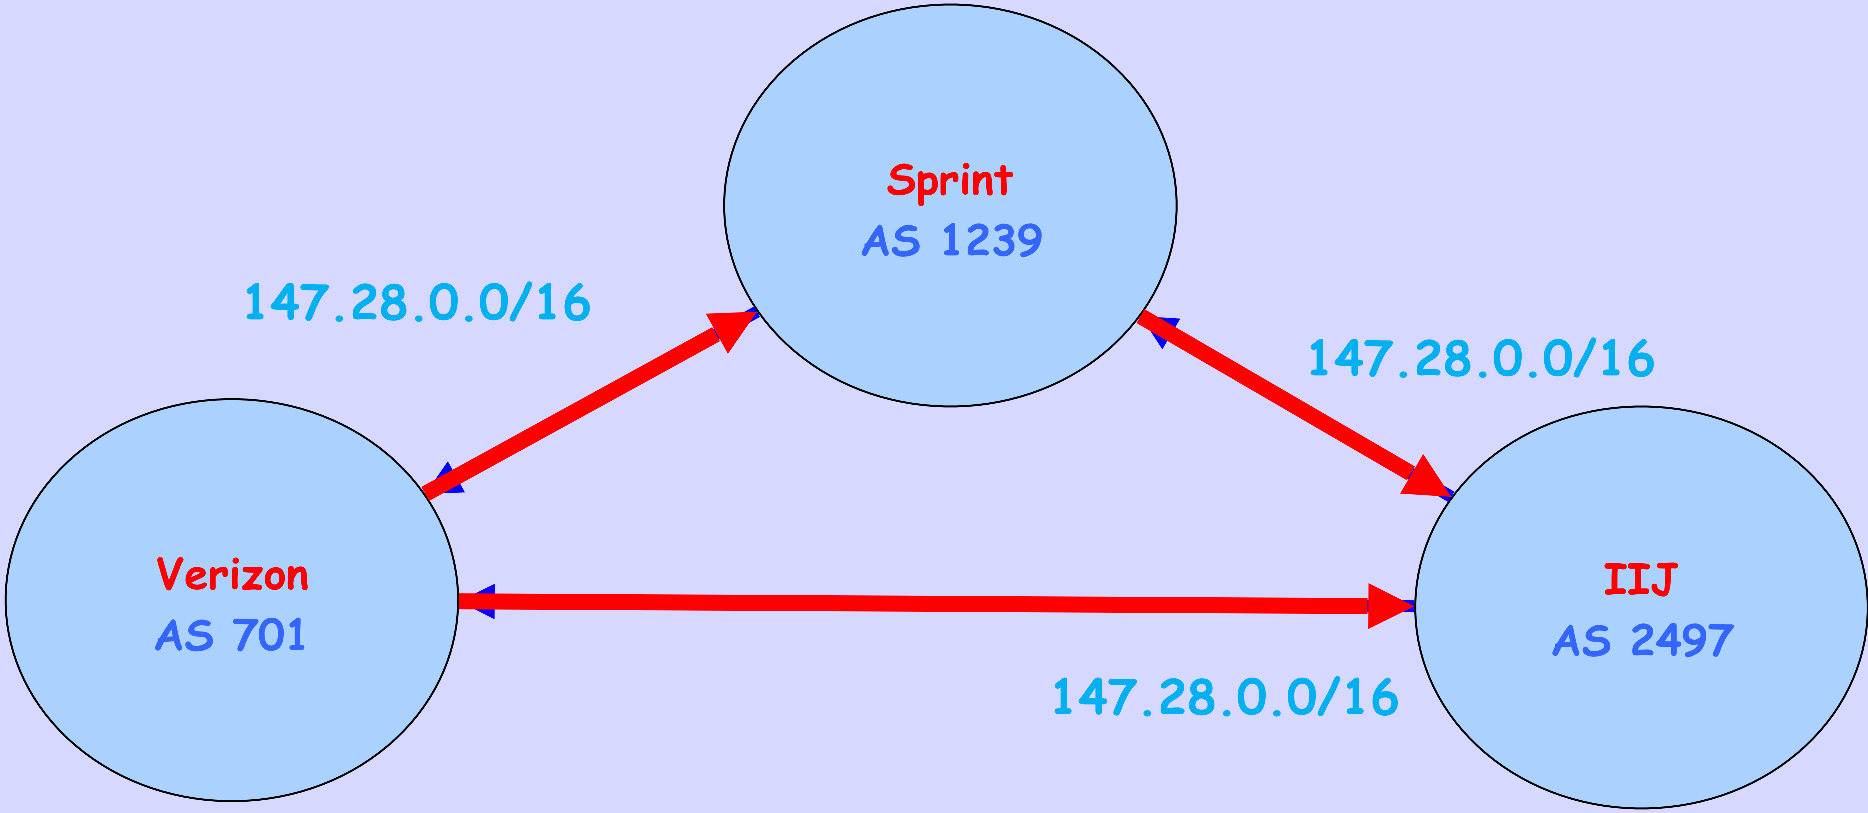

One YouTube Incident

The Plan

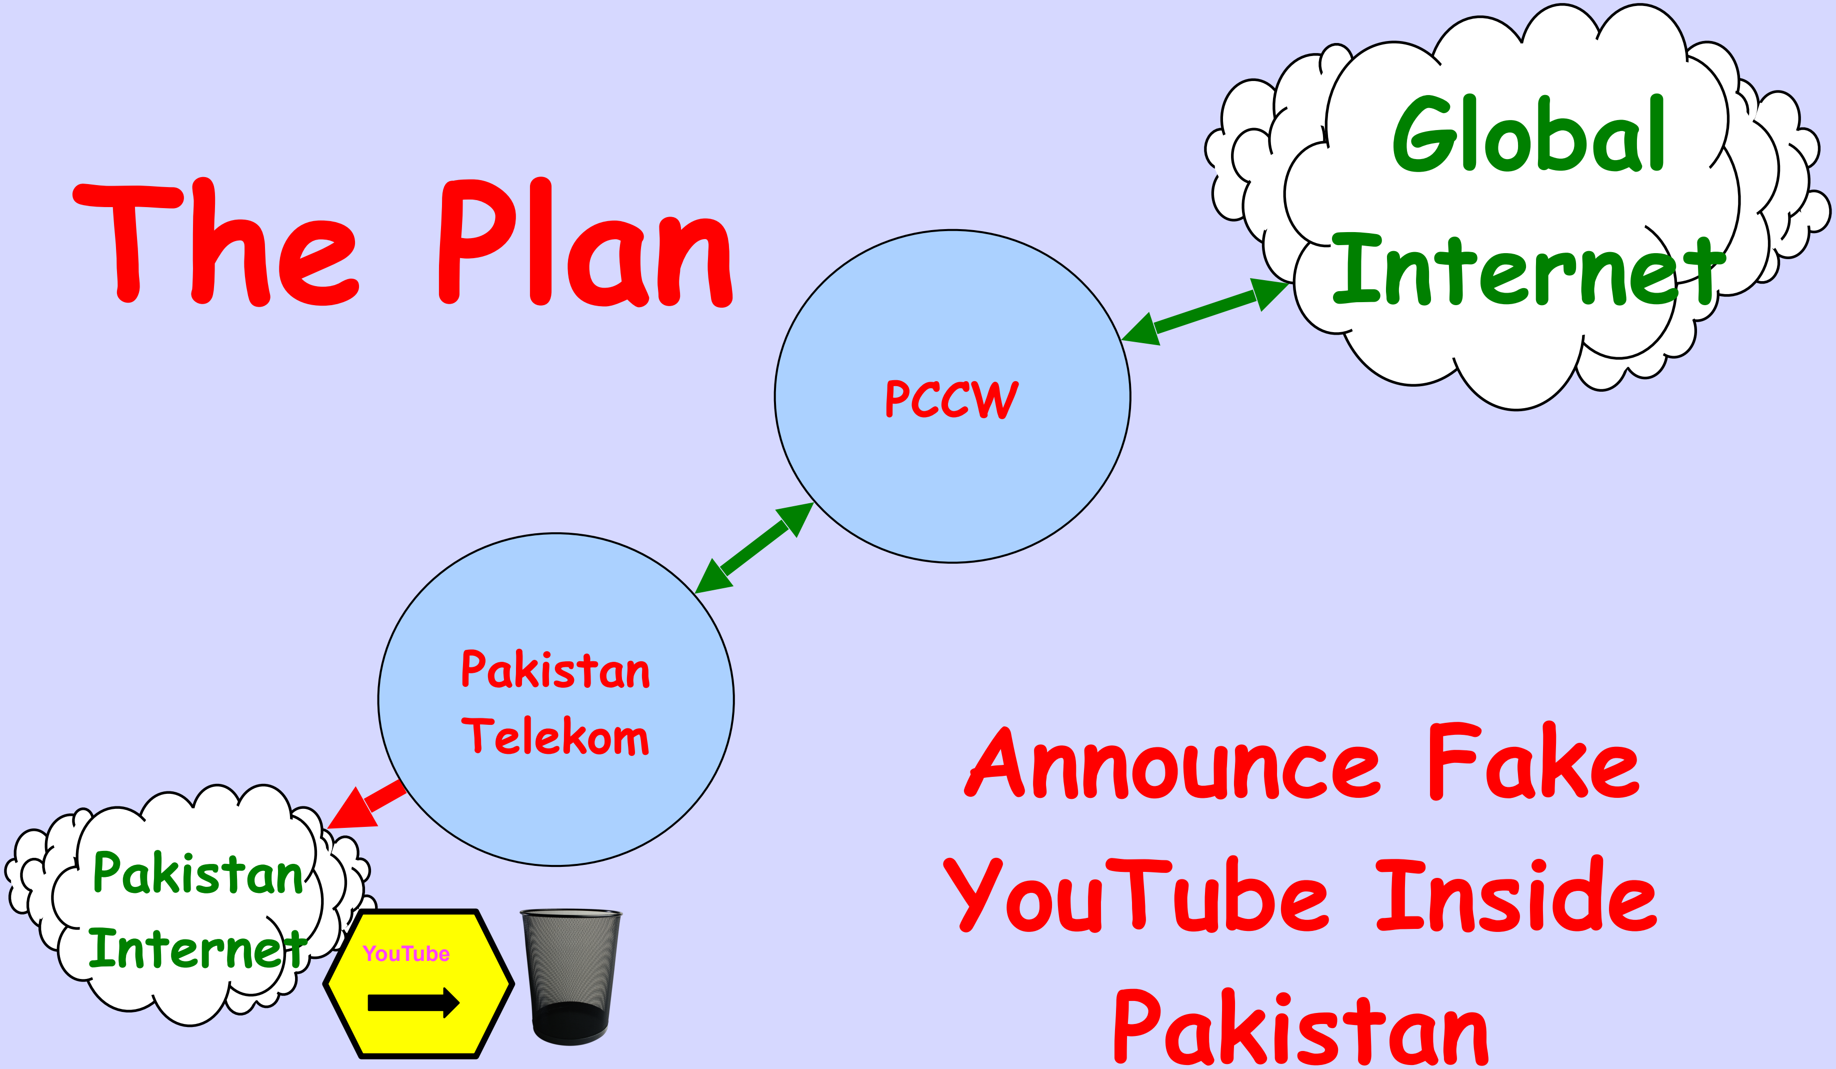

Announce Fake YouTube Inside Pakistan

What Happened

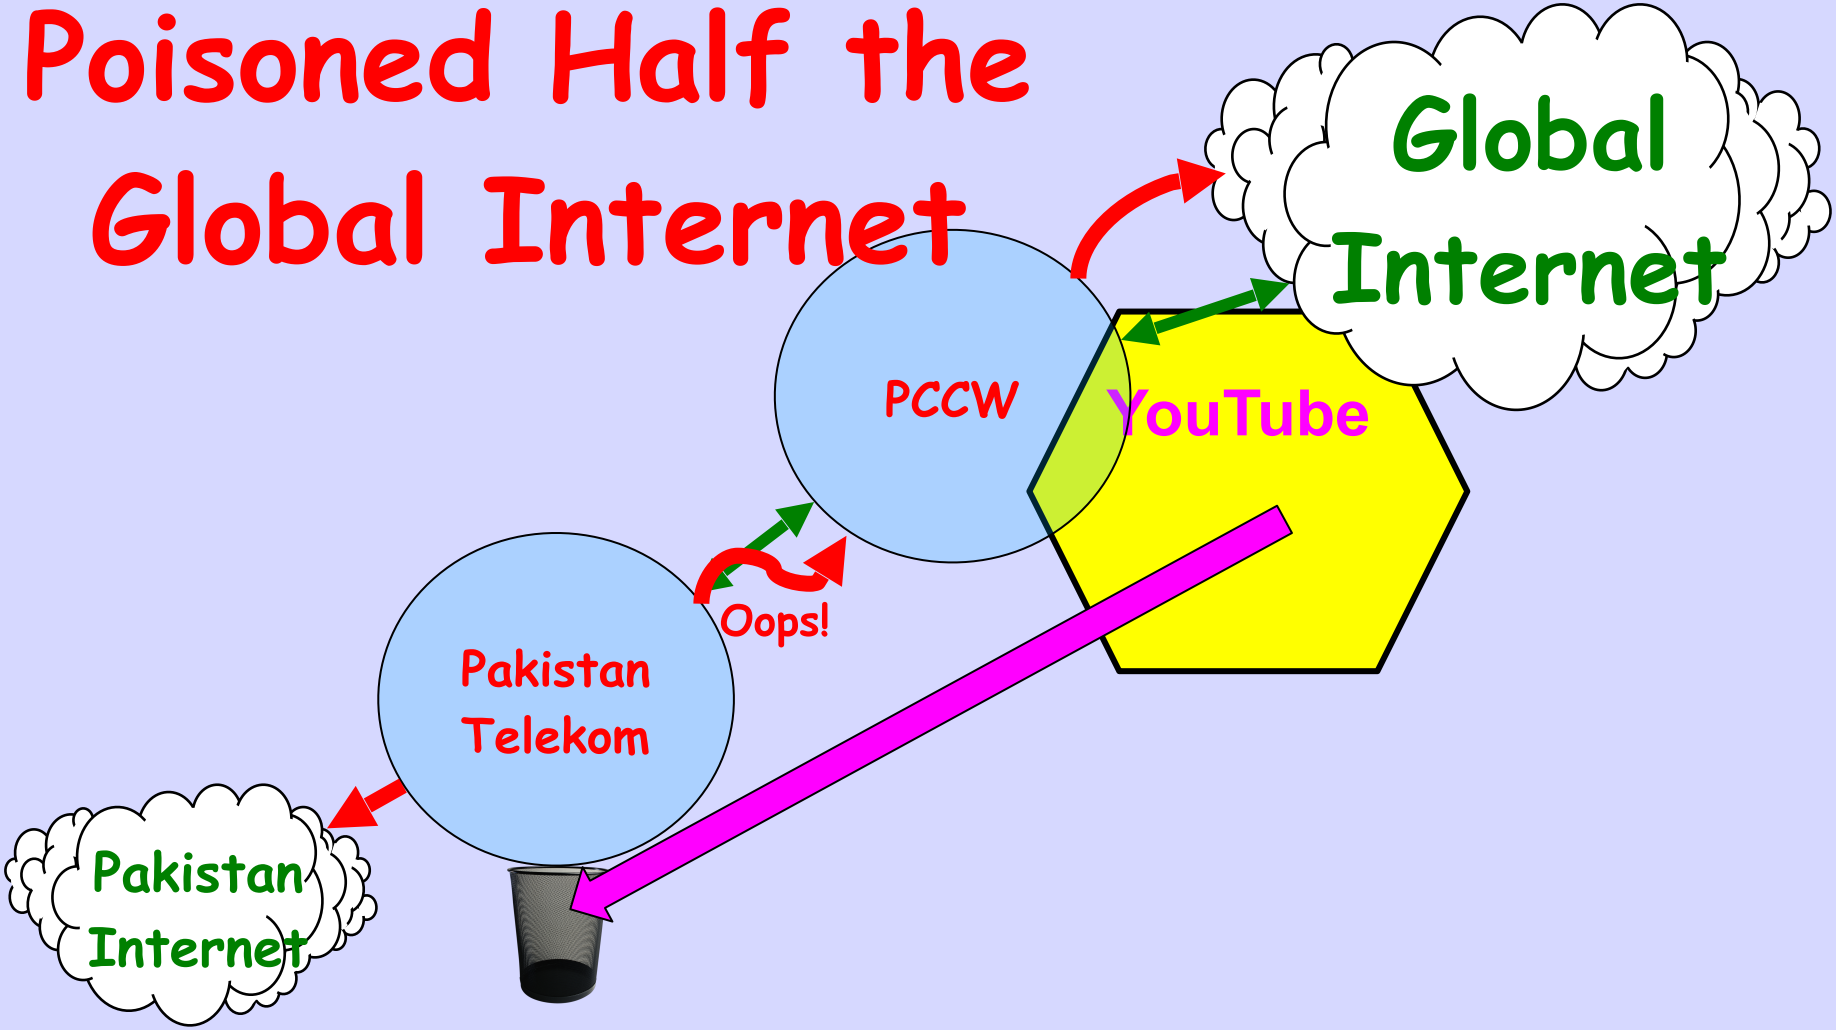

This Was Not Complex Enough

Operators Wanted Signaling on Top of Signaling

Add BGP Communities

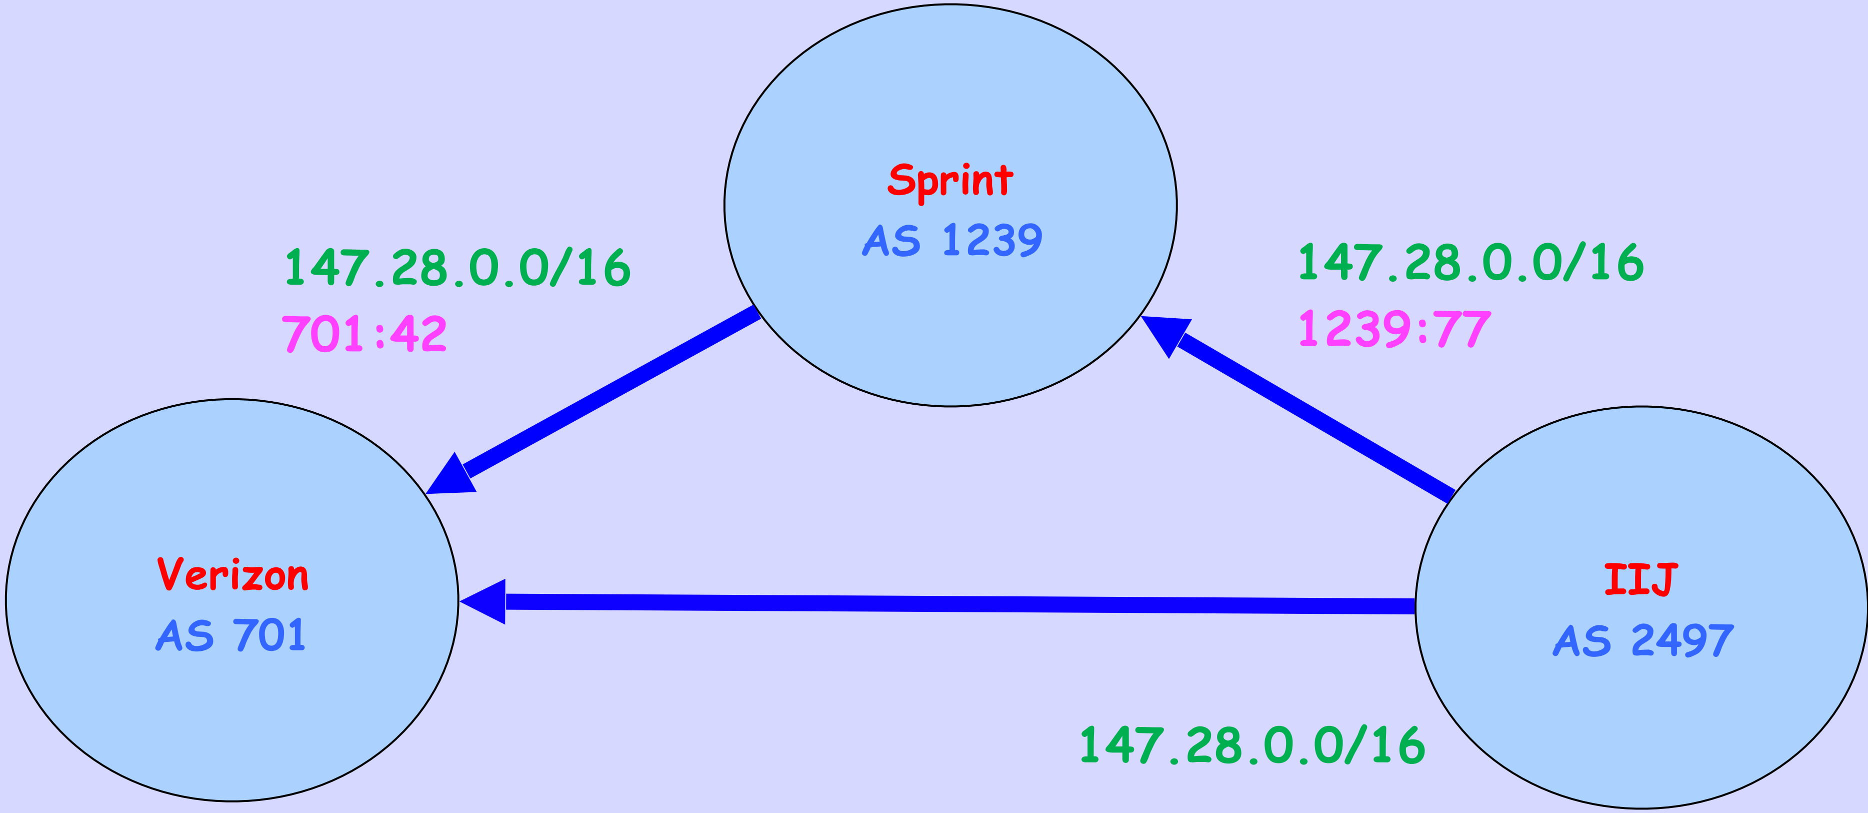

Syntax

```
AS#:number
```

But

AS#

May really be Anything

And

:number

May really Mean Anything

Undefined Semantics

We have a syntax, AS:<blarg>

But there are no formal semantics, just convention and common practice

We're putting semantics in comments

```
i = 0;  /* i = 42 */
```

Flavors, We Think

- Active
  - Path prepending
  - Modify local preference
  - Remote triggered blackholing
  - Selective announcements
- Passive
  - Location Tagging
  - RTT Tagging

And then anything a thousand kiddies have invented

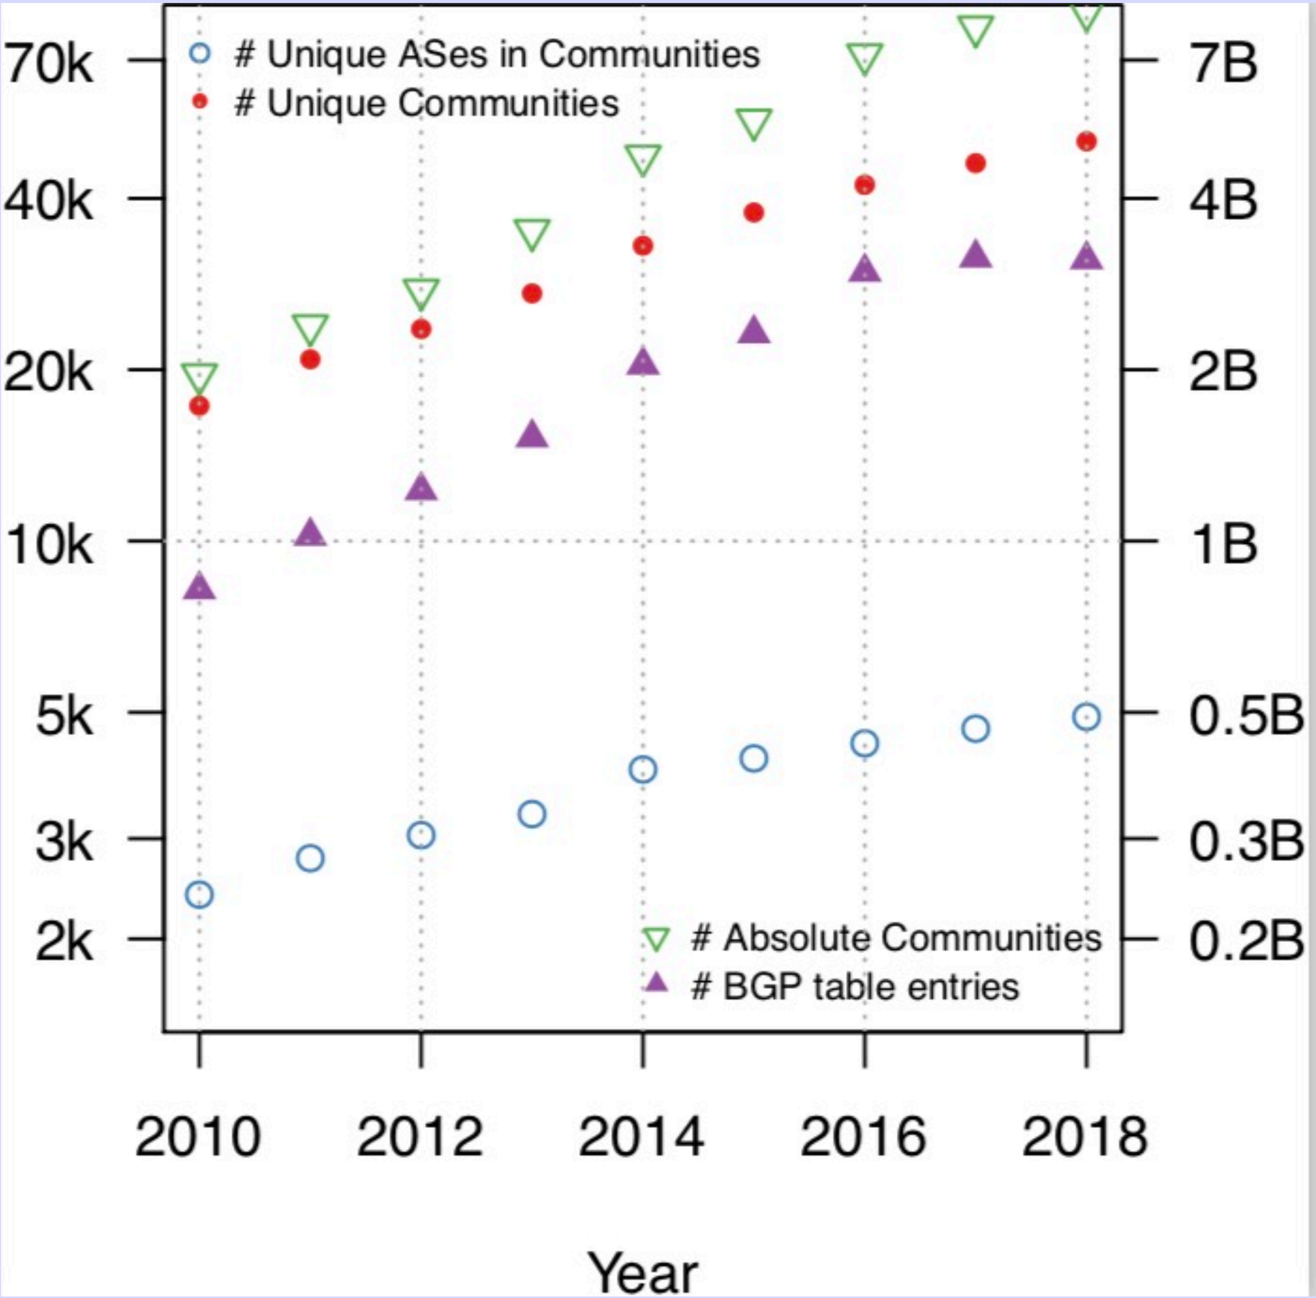
○ # Unique ASes in Communities
● # Unique Communities
▽ # Absolute Communities
▲ # BGP table entries
70k
40k
20k
10k
5k
3k
2k
7B
4B
2B
1B
0.5B
0.3B
0.2B
2010
2012
2014
2016
2018
Year

Propagation

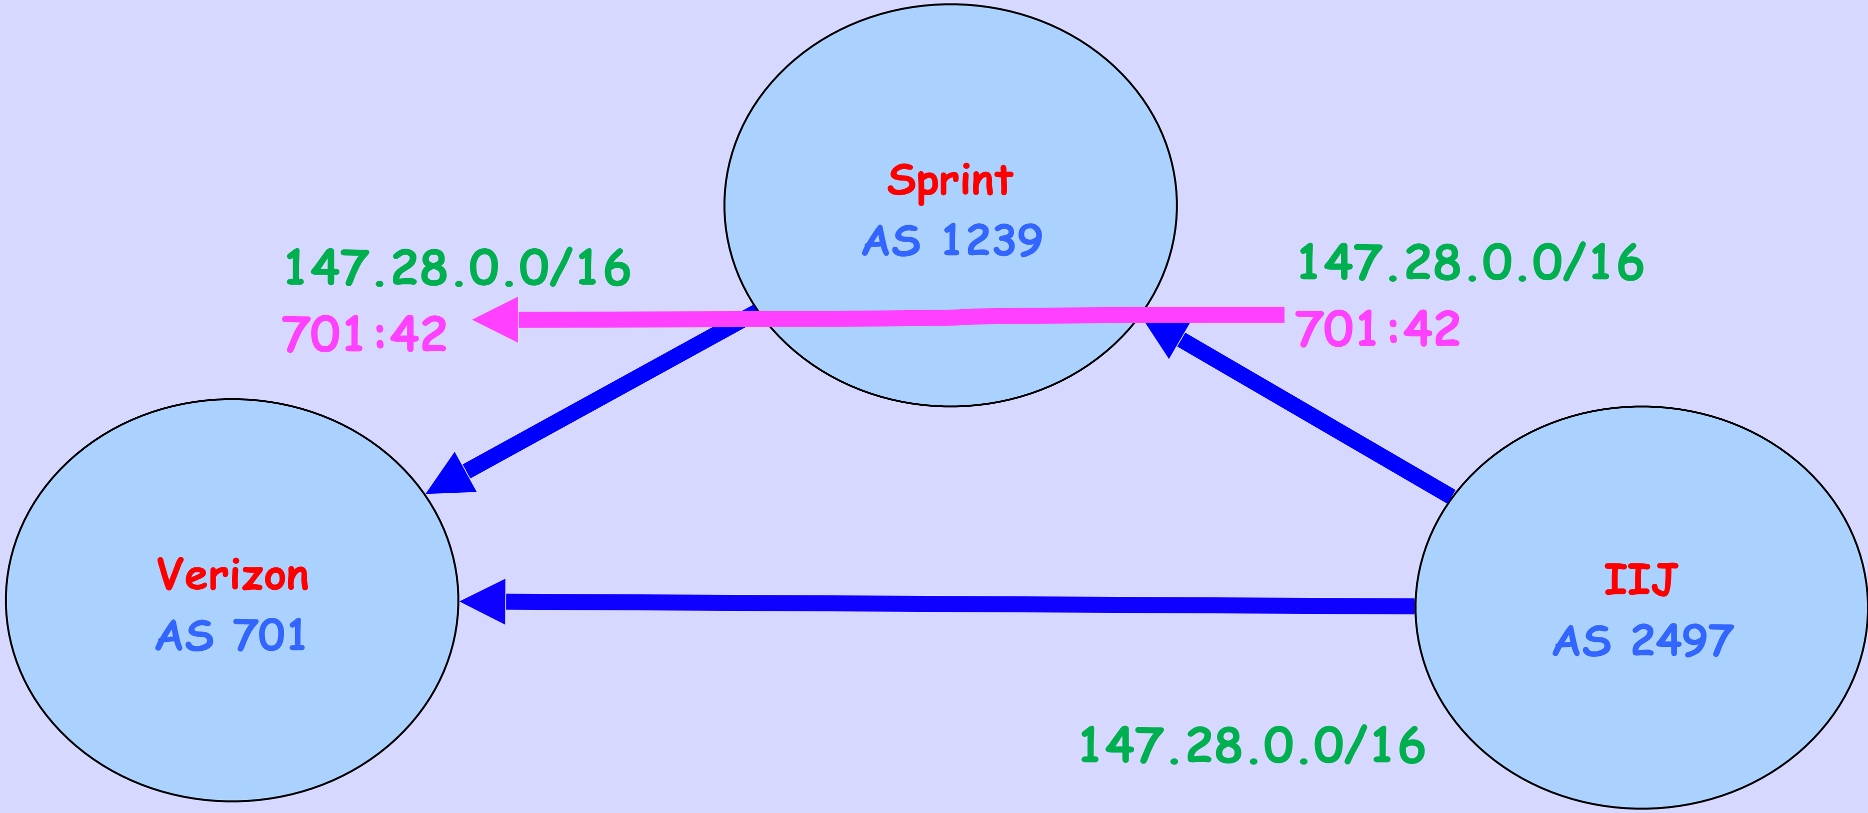

Propagation

- RFC 1997: Communities are a transitive optional attribute
- RFC 7454: Scrub own, forward foreign communities
- So many people do not expect them to propagate that widely
- I, for one, did not

Only 14% of Transit ASs propagate communities

(2.2k of 15.5k)

Surprise!

- 14% seems small, but the AS graph is highly connected
- More than 50% of communities traverse more than four ASes
- 10% of communities have a hop count of more than six ASes
- Longest community propagation observed: through 11 ASes

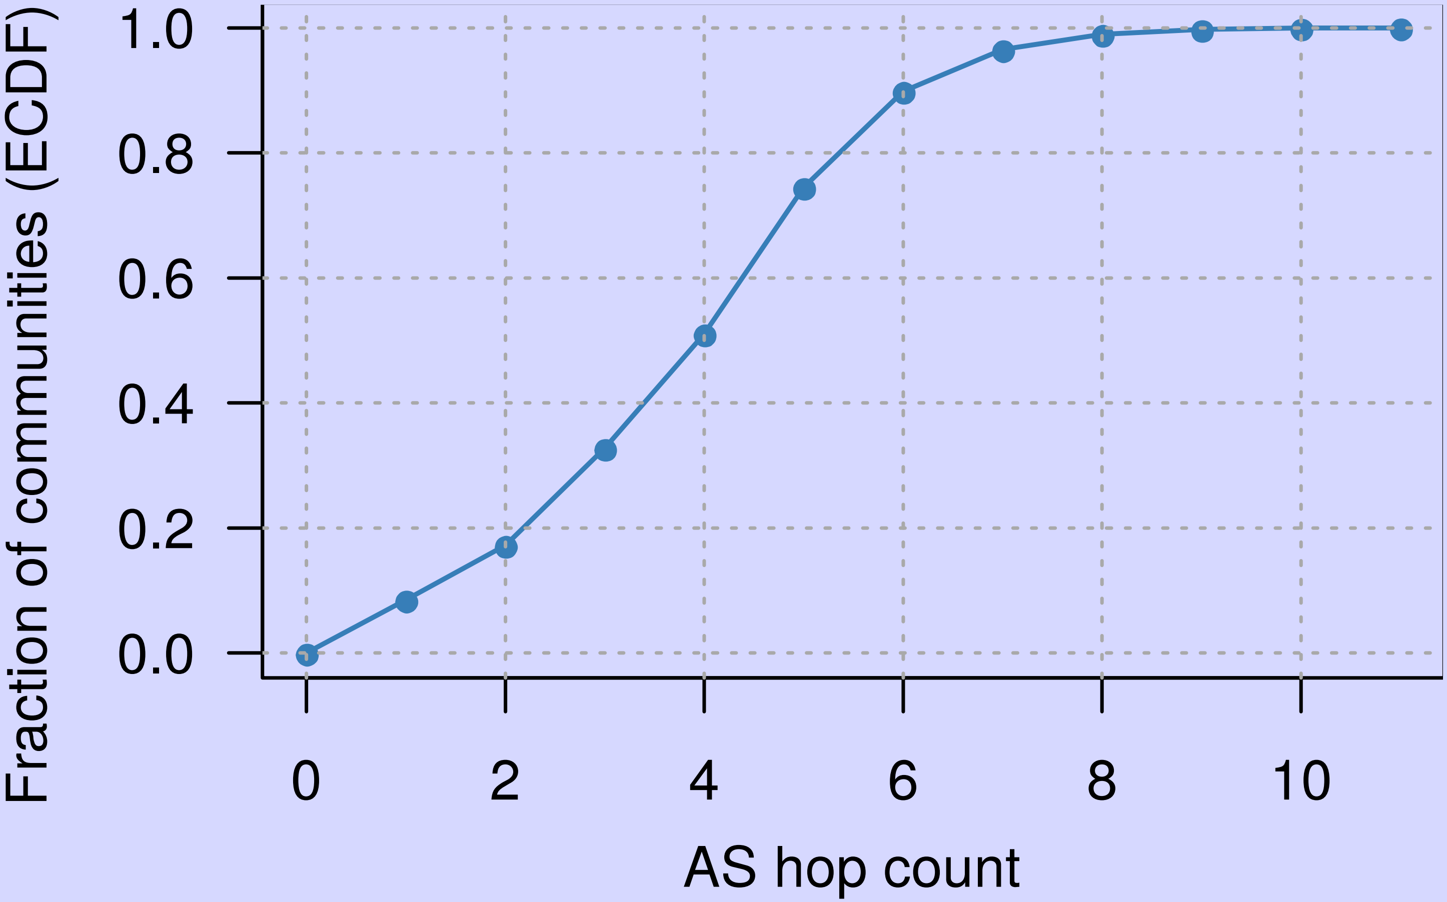
Fraction of communities (ECDF)
1.0
0.8
0.6
0.4
0.2
0.0
0
2
4
6
8
10
AS hop count

On/Off Path

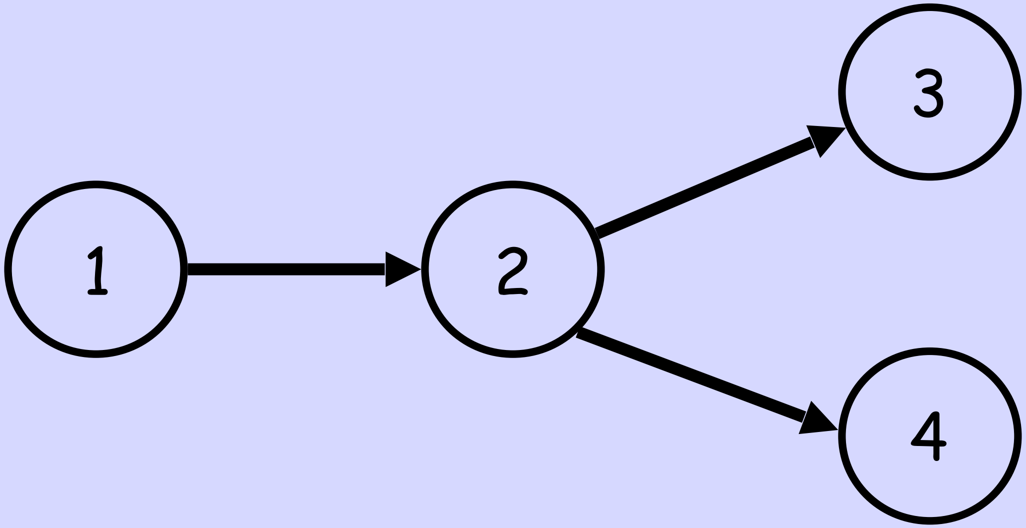

On/Off Path

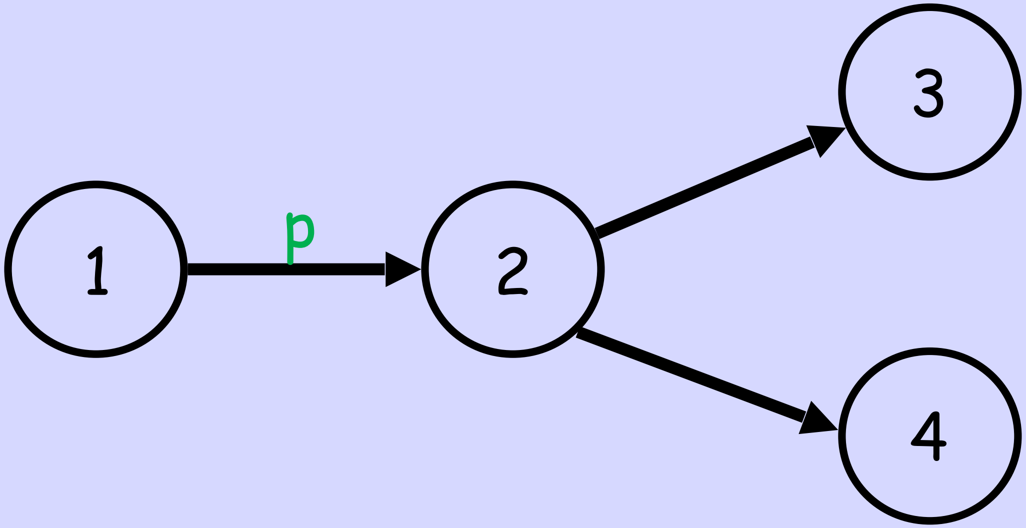

On/Off Path

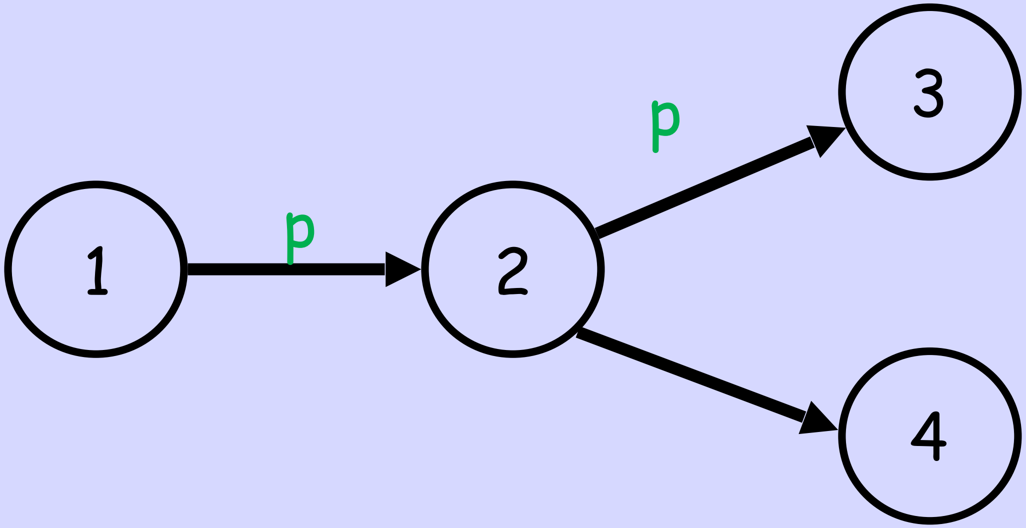

On/Off Path

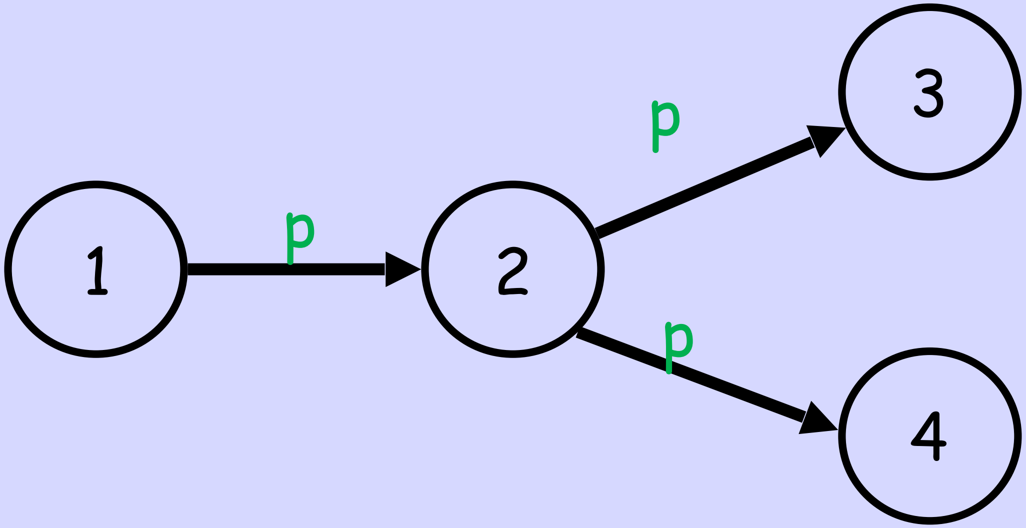

On/Off Path

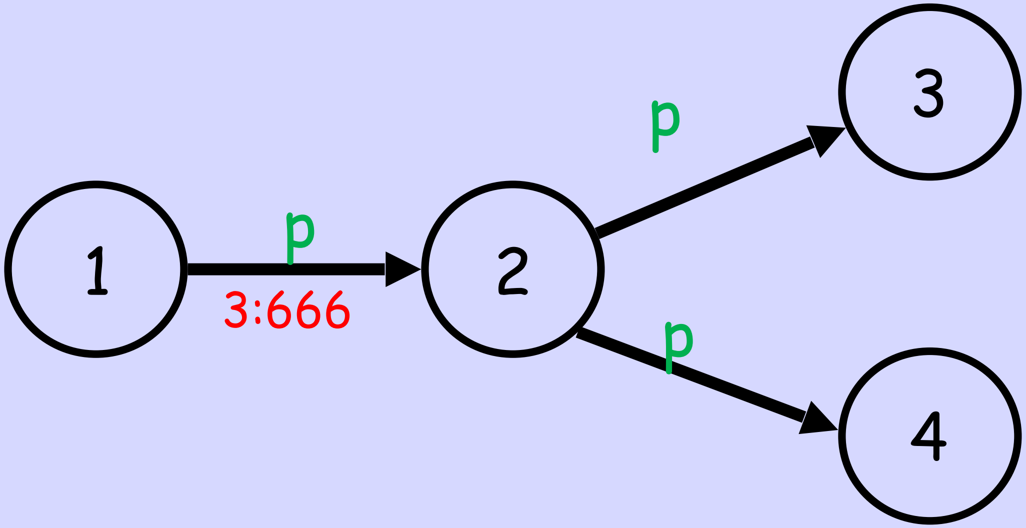

On/Off Path

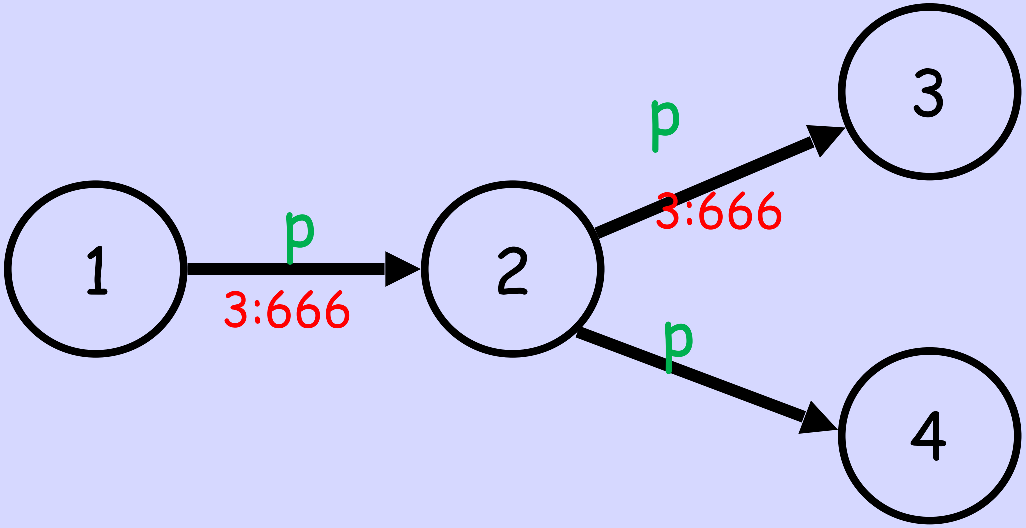

On/Off Path

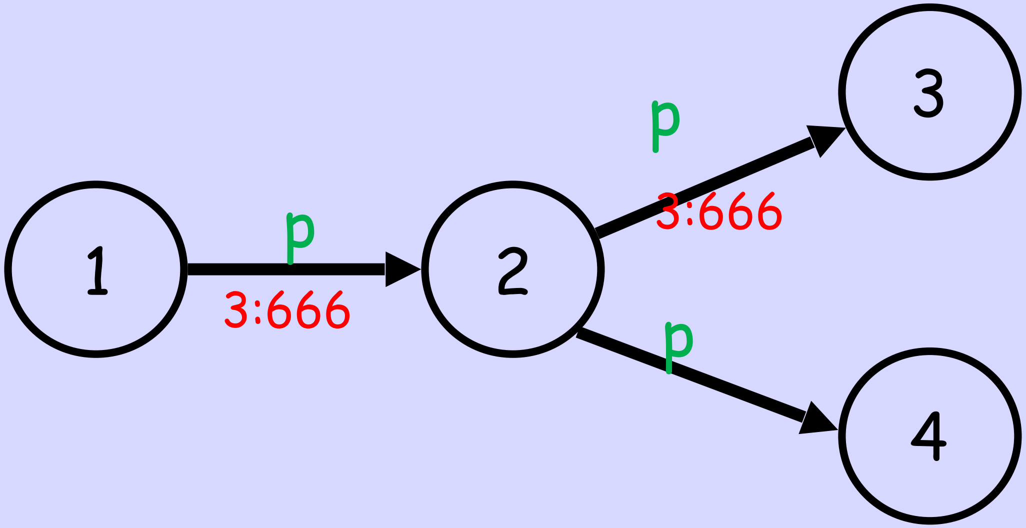

2 and 3 are On Path

On/Off Path

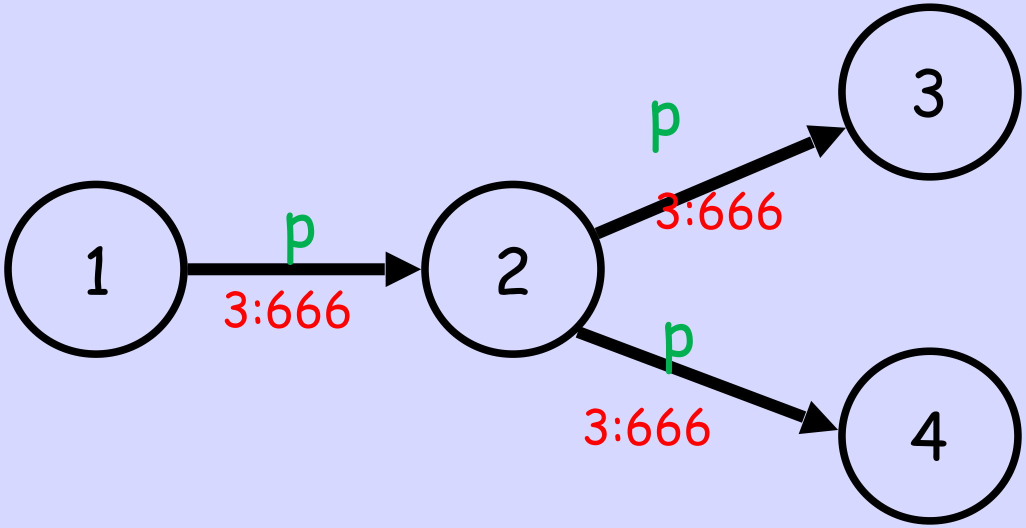

2 and 3 are On Path

On/Off Path

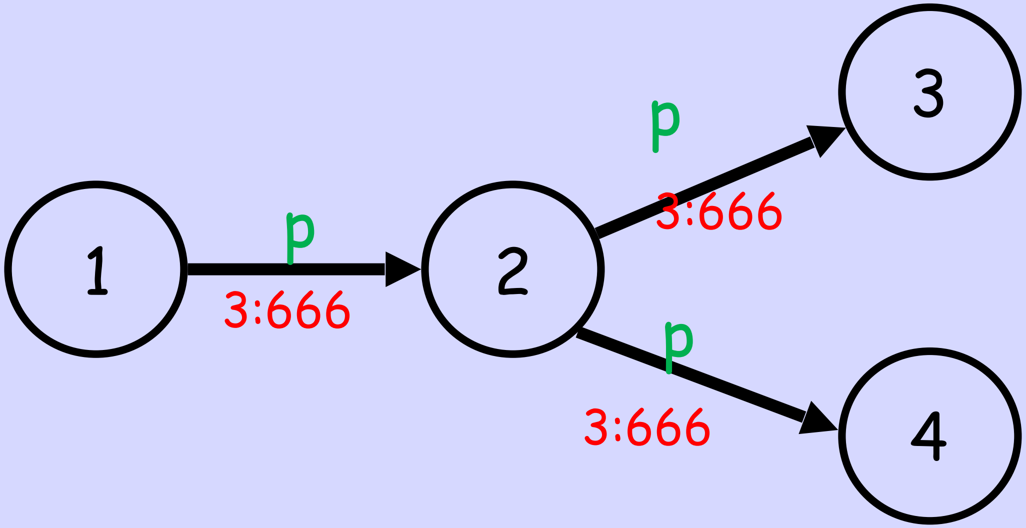

2 and 3 are On Path

4 is Off Path

Observed Communities

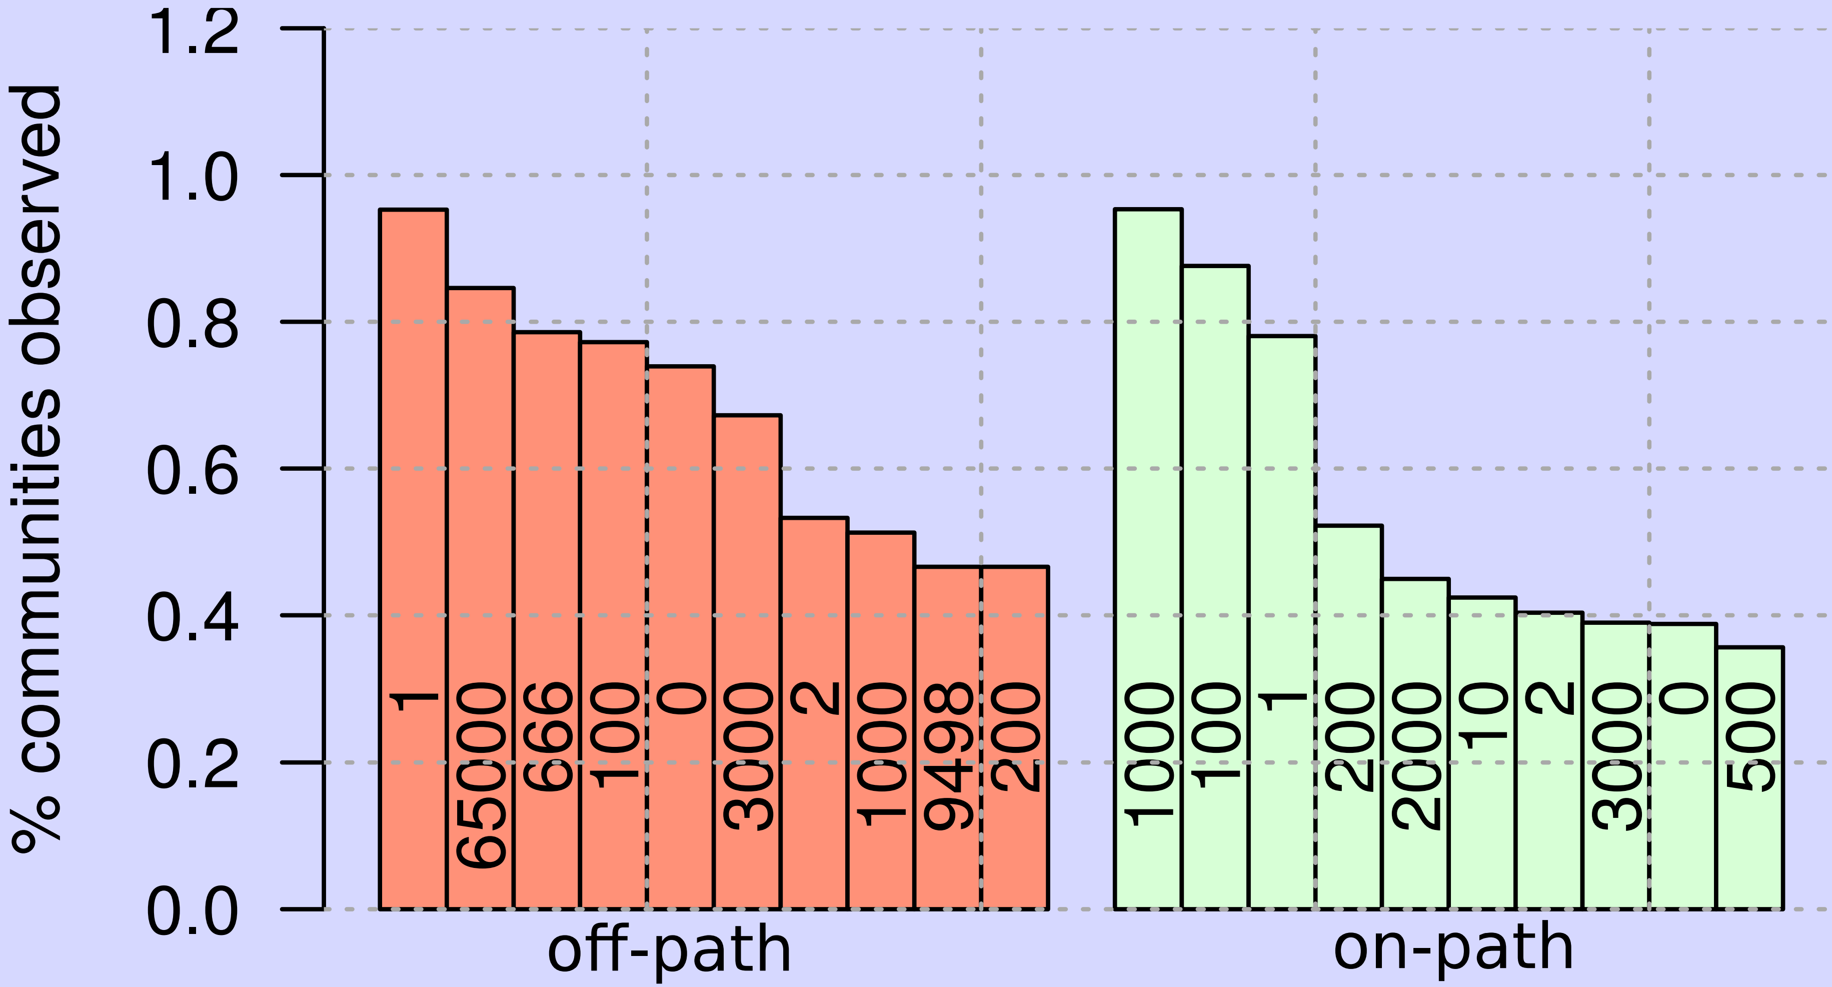

And We Have No Idea What Almost All of Them Mean

The Internet is an Experimental Hack

So Let's Break Things!

Method to our Madness

- All experiments first tested in Lab
- Impacts were estimated
- Validated on the Internet, with operators' consent, e.g. for hijacks

RTBH

One of the Very Few Defined Communities

RTBH

Remotely Triggered Black Hole Community

Target-AS:666

Attached to a Prefix

A DoS Defense

Signaling that Traffic to a Prefix be Dropped

DoS Attack

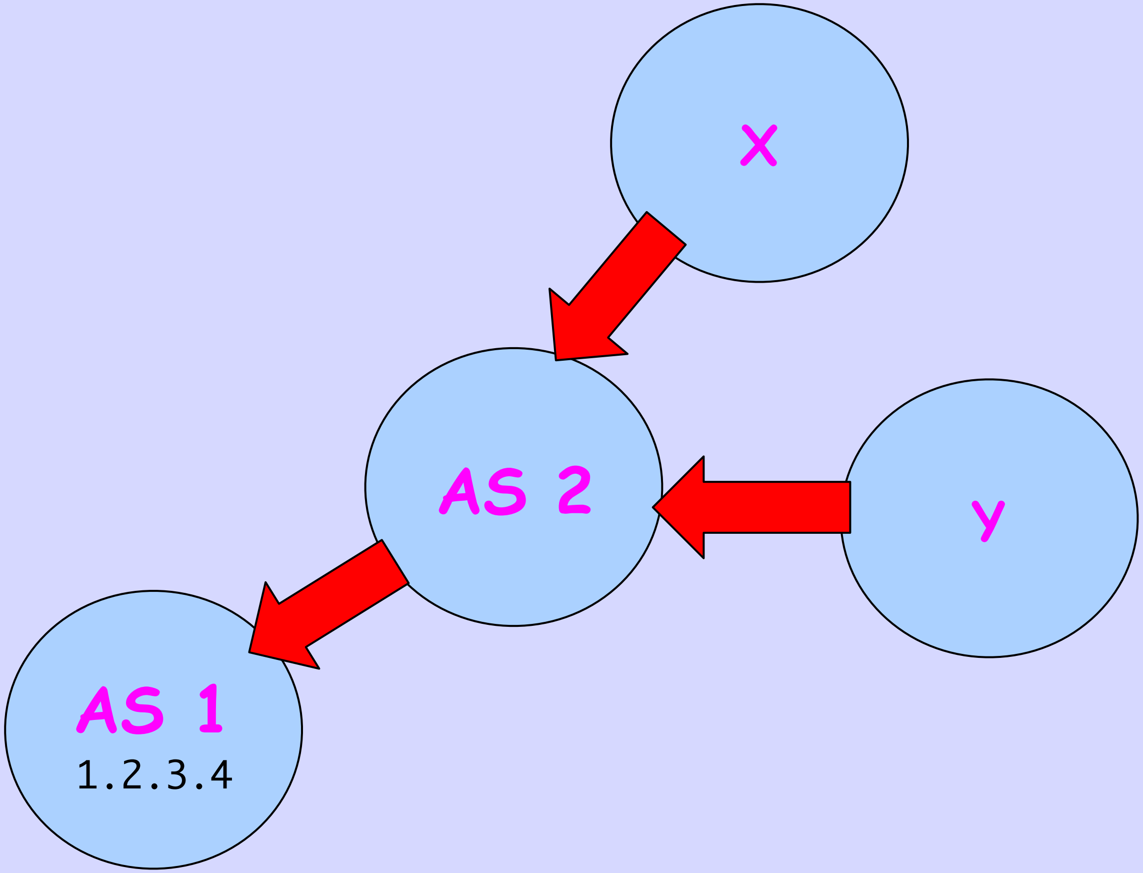

Creative Commons: Attribution & Share Alike

Ask AS 2 to Black Hole

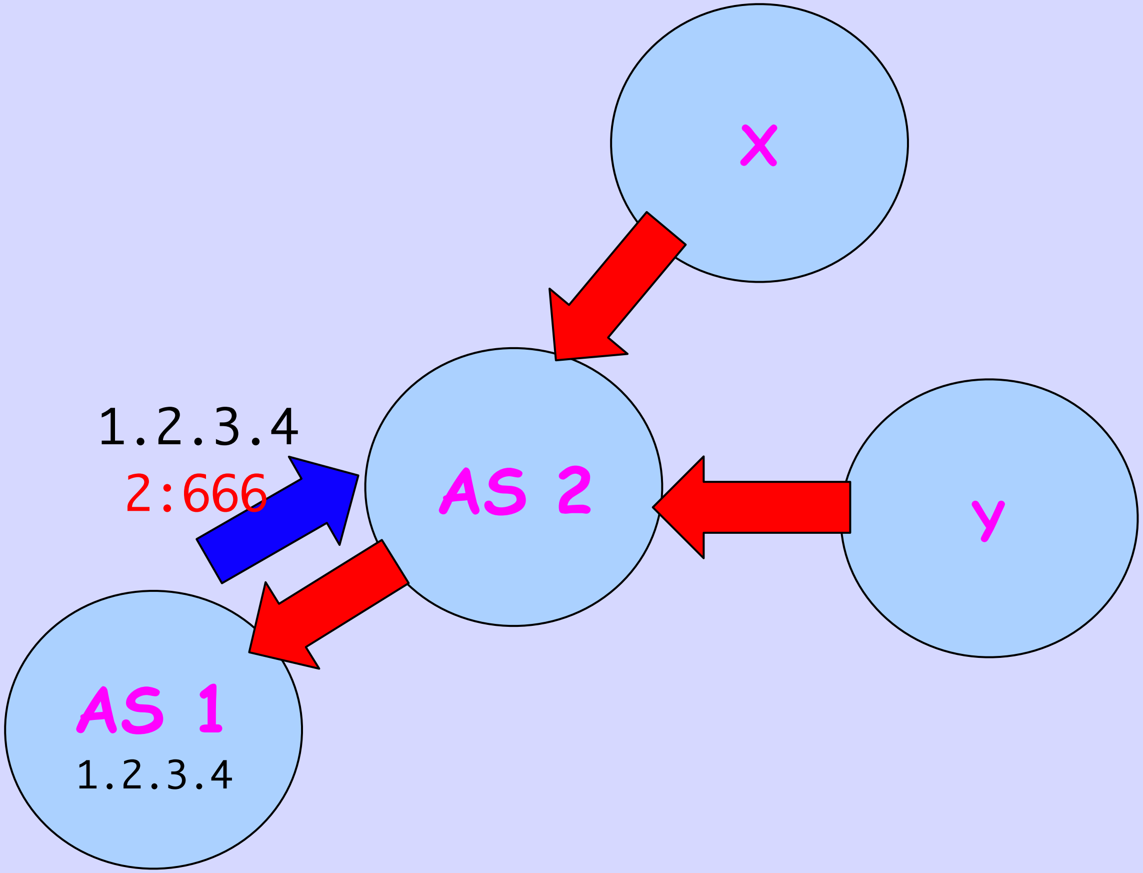

Traffic Dropped

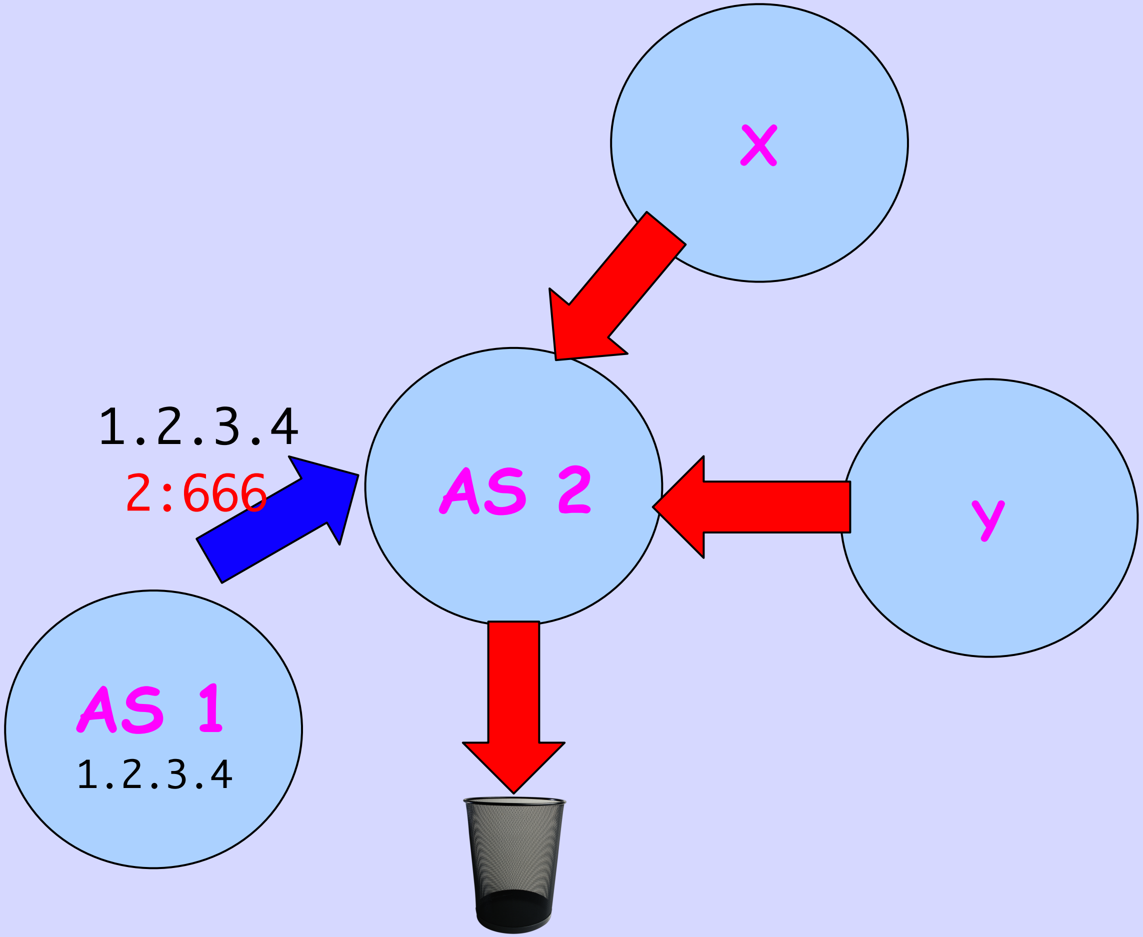

Safeguards, in Theory

- Provider should check customer prefix before accepting RTBH
- Customer may only blackhole own prefixes
- Different policies for Customers/Peers
- On receiving RTBH, do not propagate

Which Looks Very Cool

Except it is an Attack Vector

The Attack

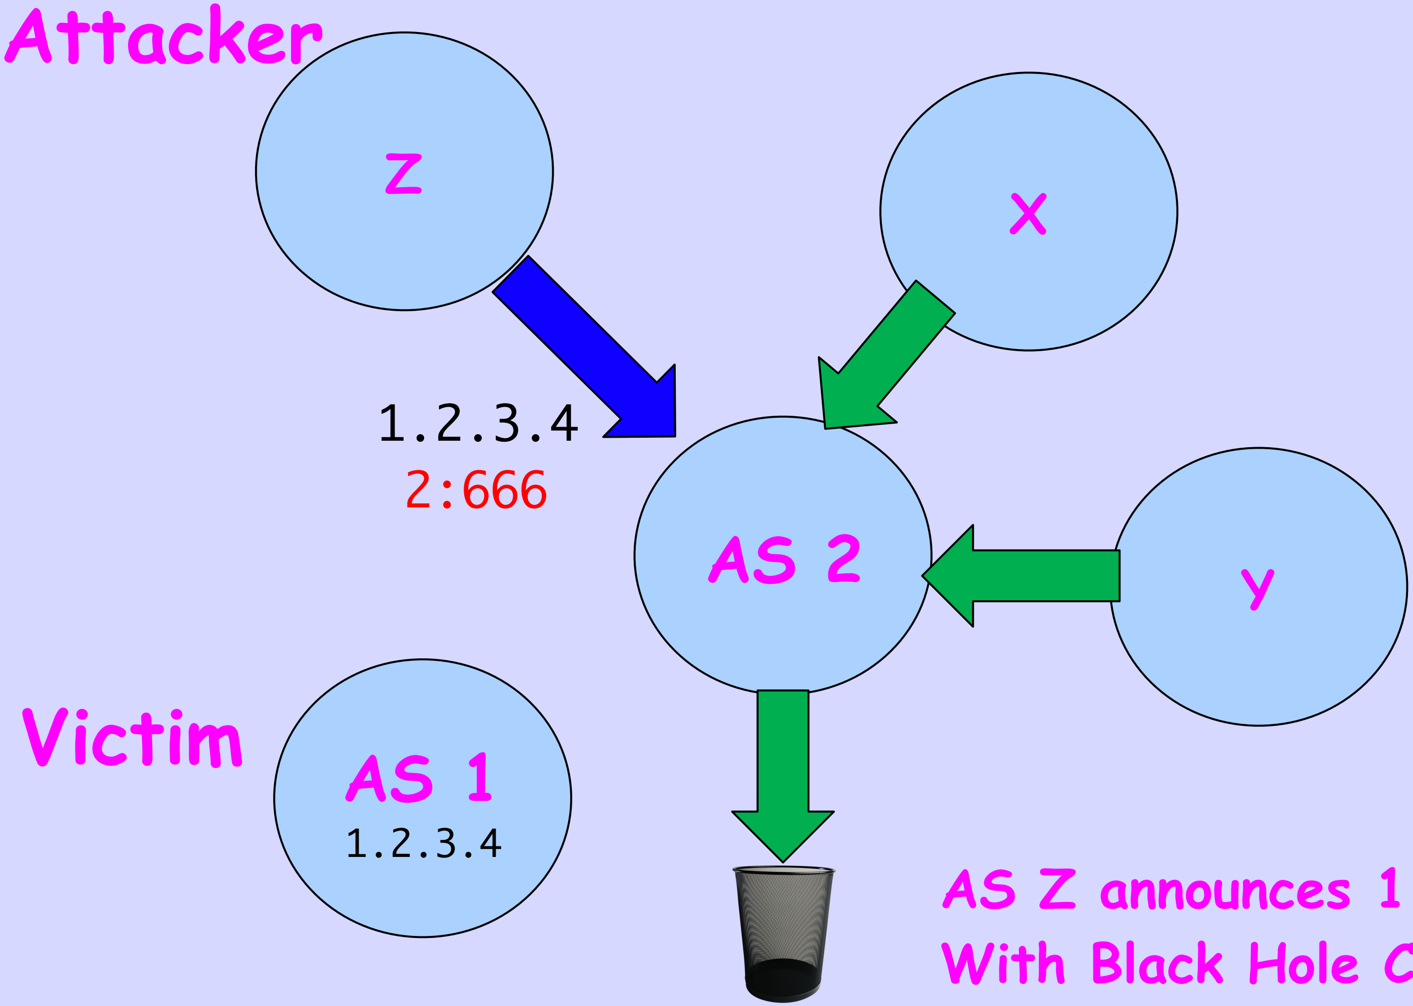

AS Z announces 1.2.3.4 to AS 2
With Black Hole Community
Good Traffic to AS 1 is Dropped

The Attack Works Well

- Works from a distance and is hard to spot
- Triggering RTBH is possible for attackers because, e.g.,:
  - BH prefix is more specific, thus accepted via exception
  - Providers check BH community before prefix filters (bug in NANOG recipe)
  - No validation for origin of community is possible

Off-Path Attacks

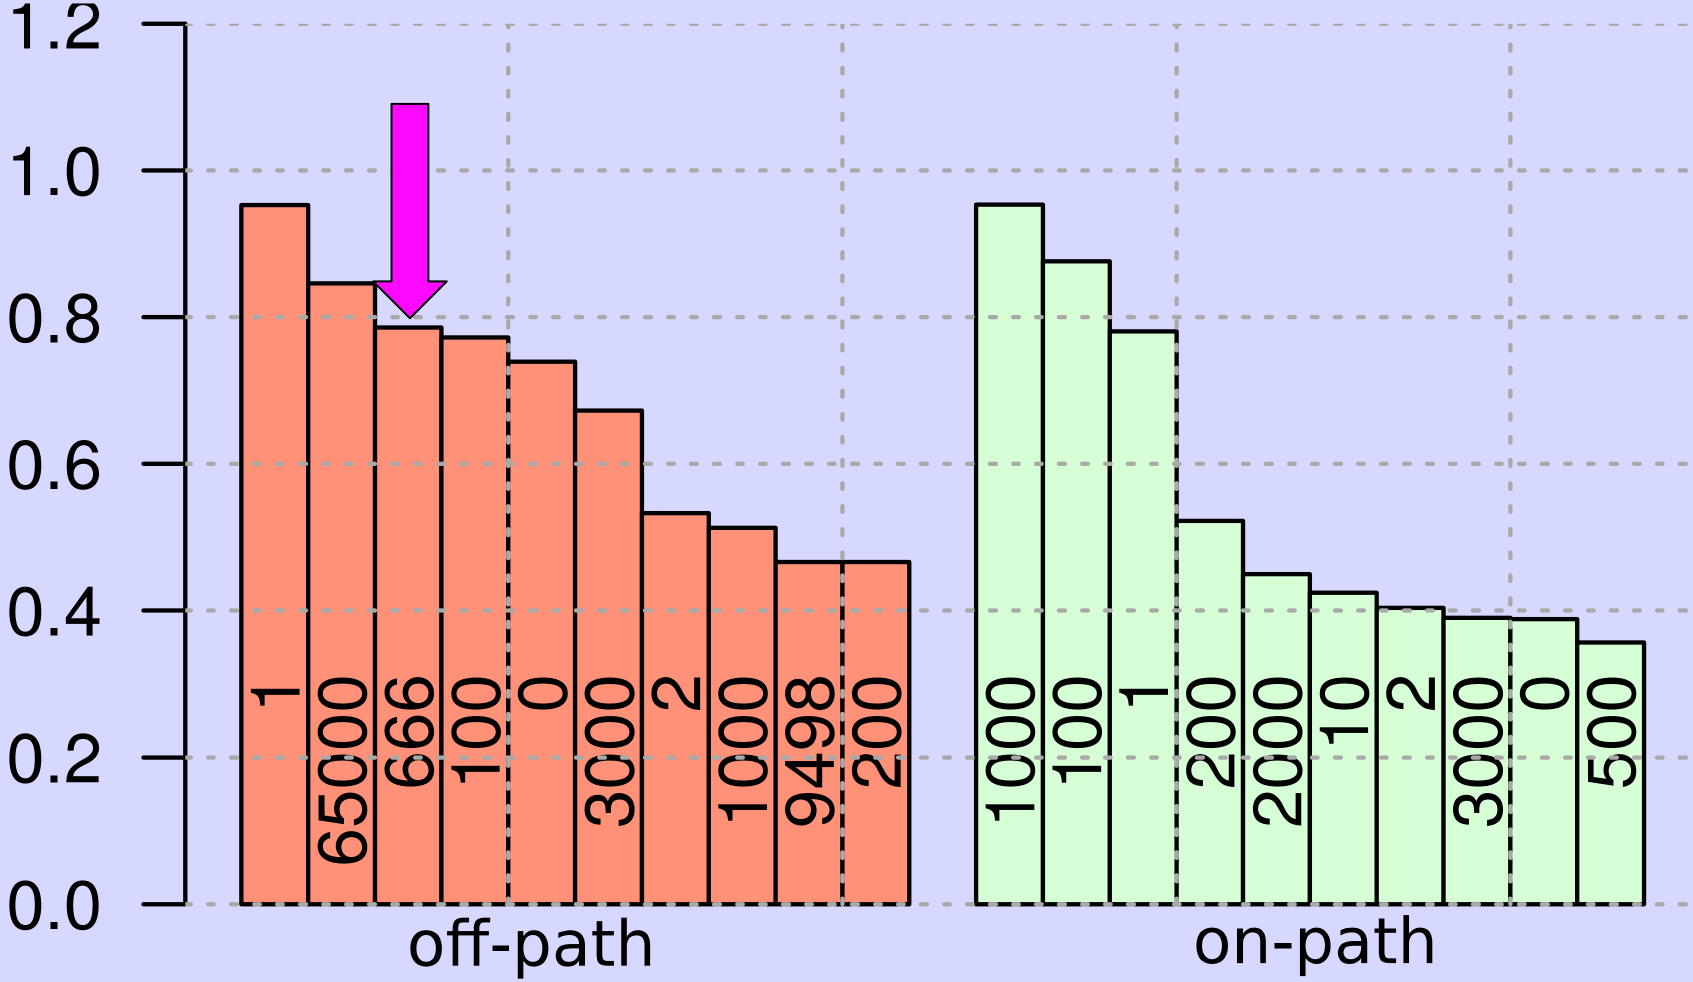

Traffic Steering

p 4 3 2 1

Traffic Steering

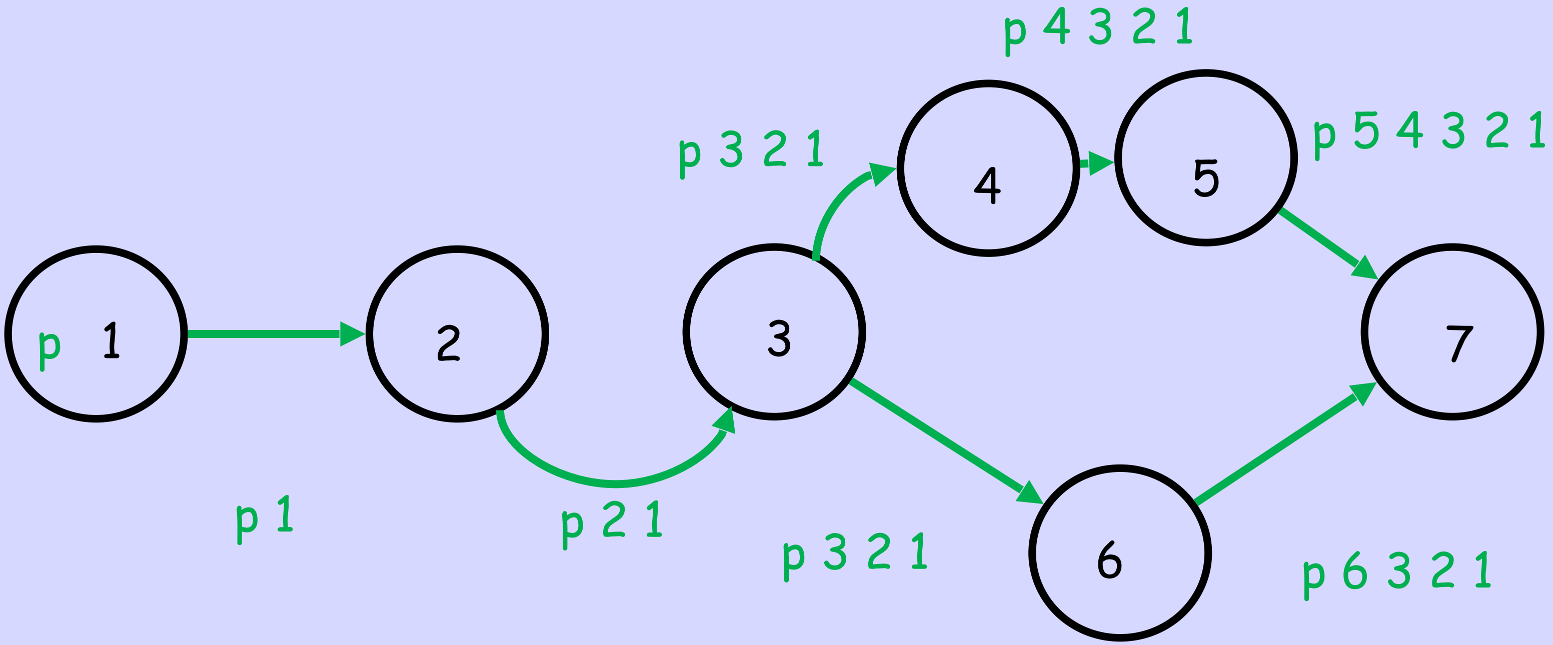

Traffic Steering

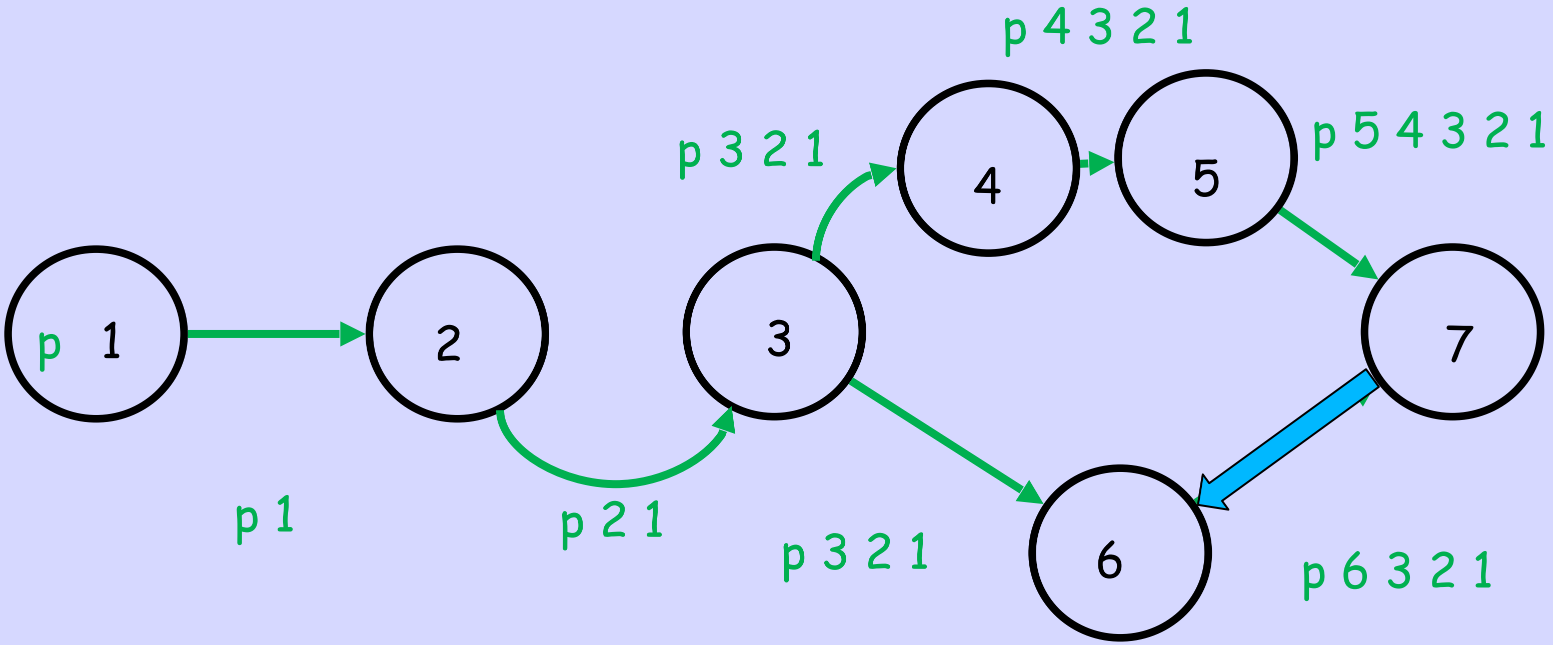

Traffic Steering

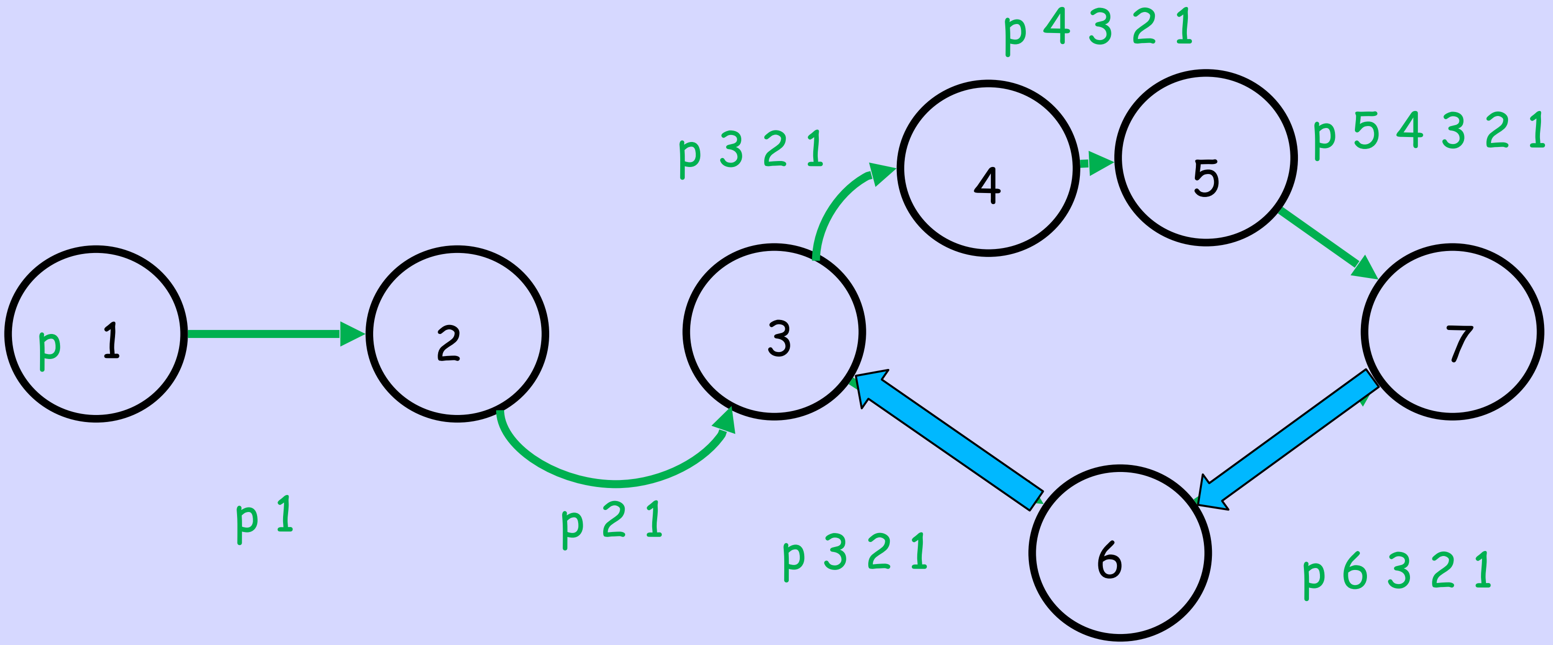

Traffic Steering

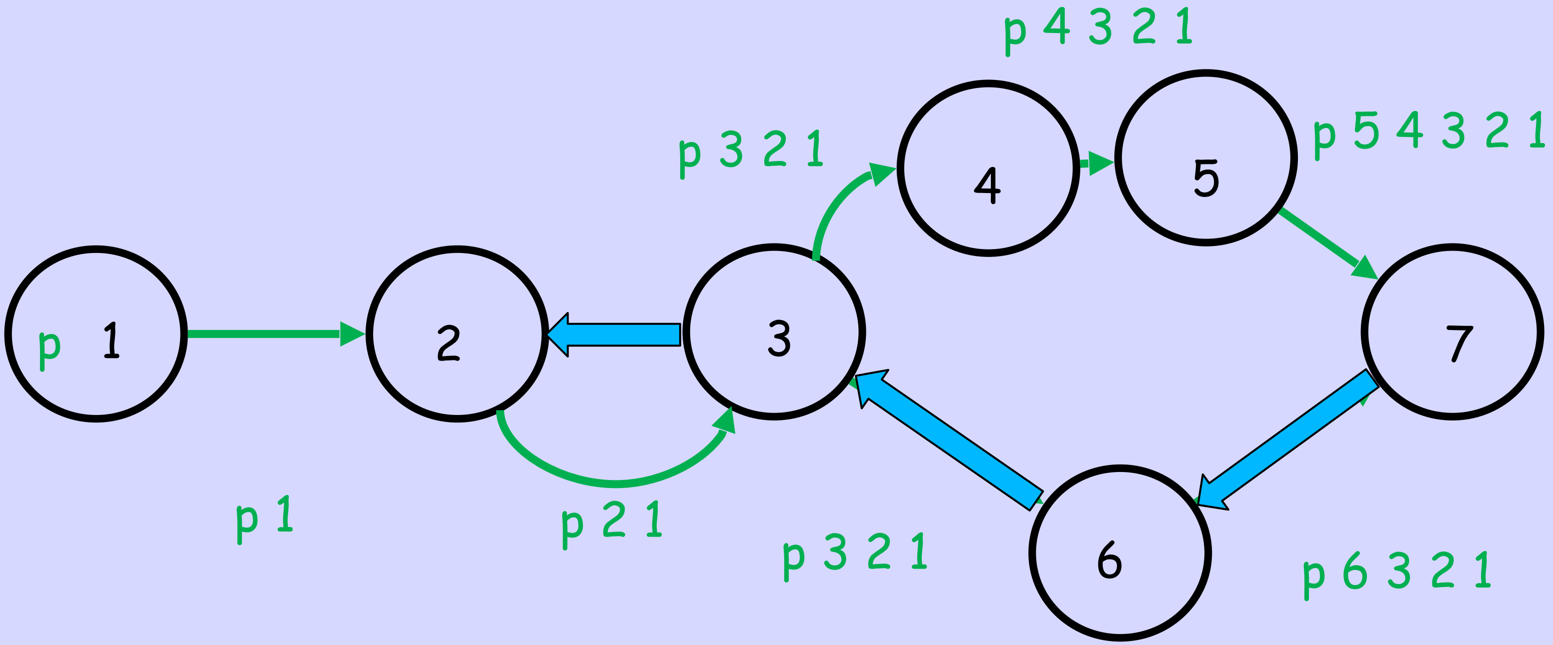

Traffic Steering

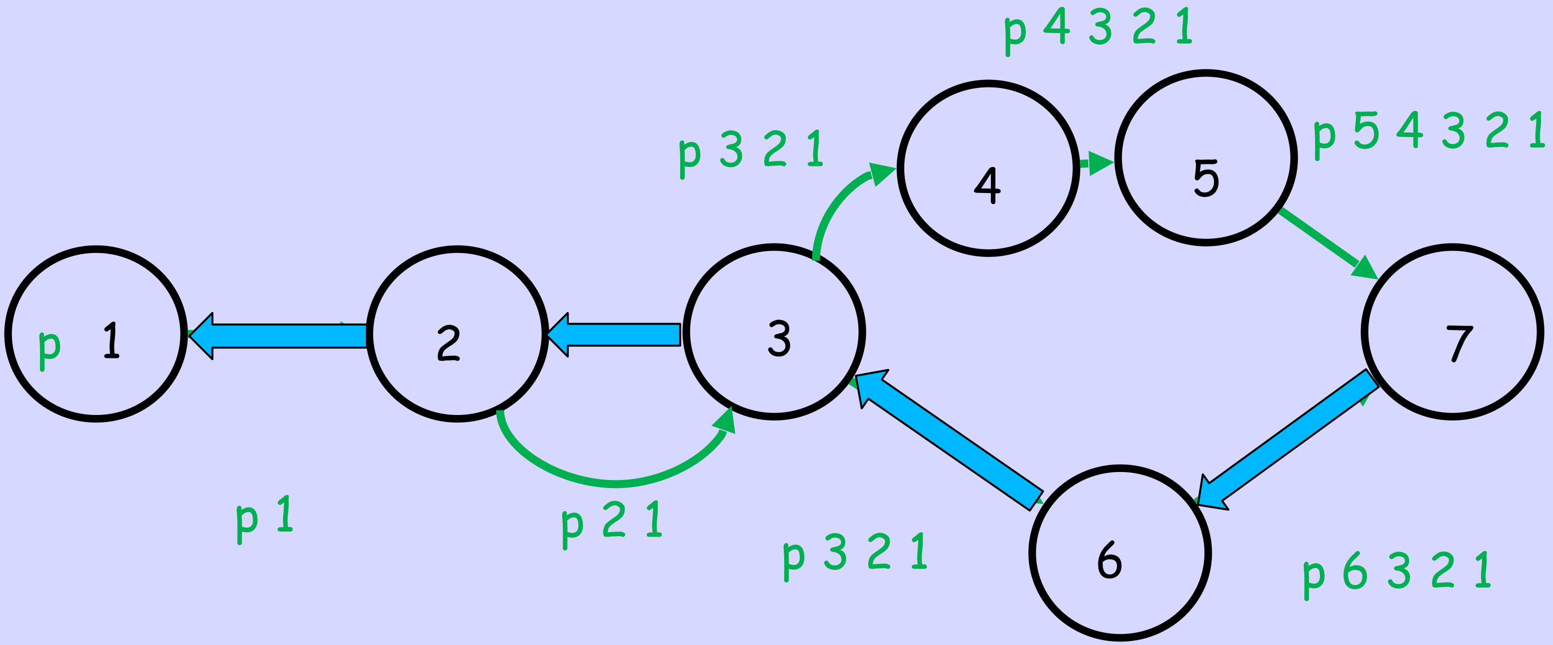

Traffic Steering

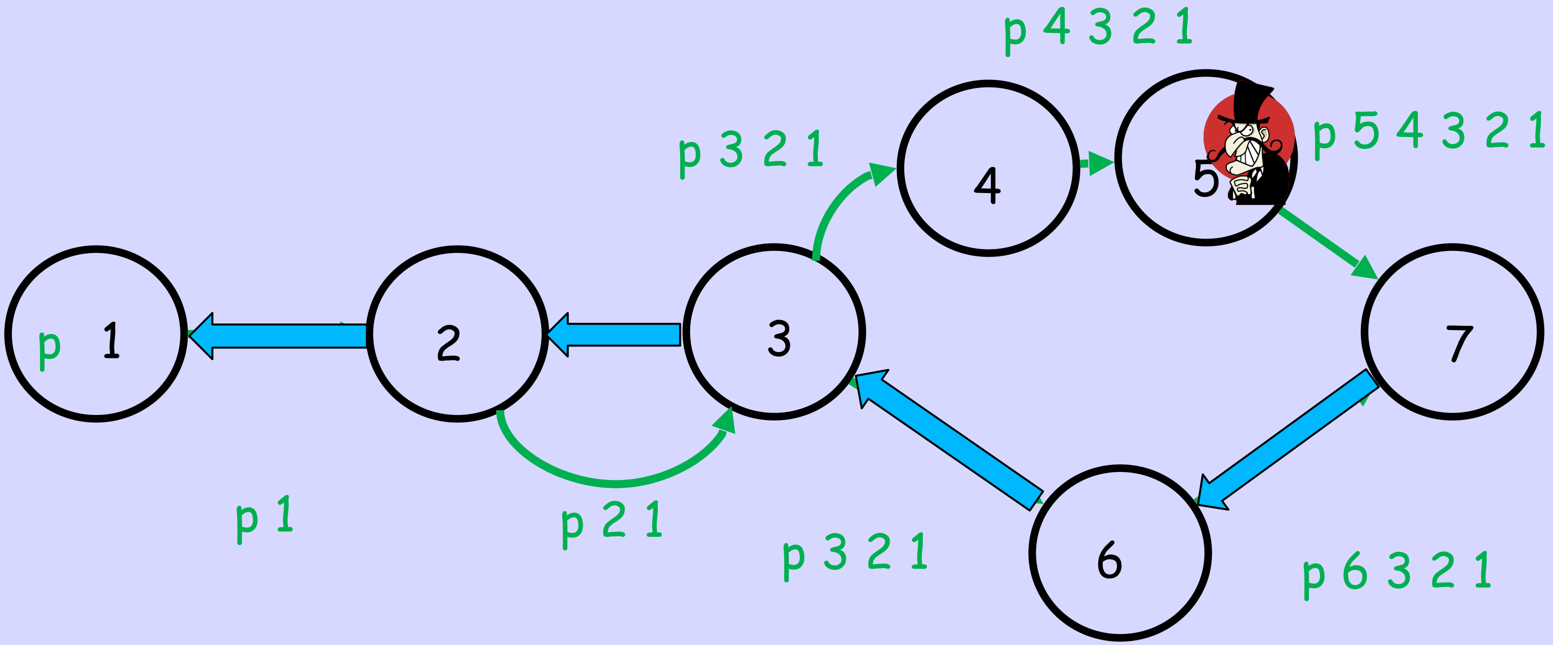

Traffic Steering

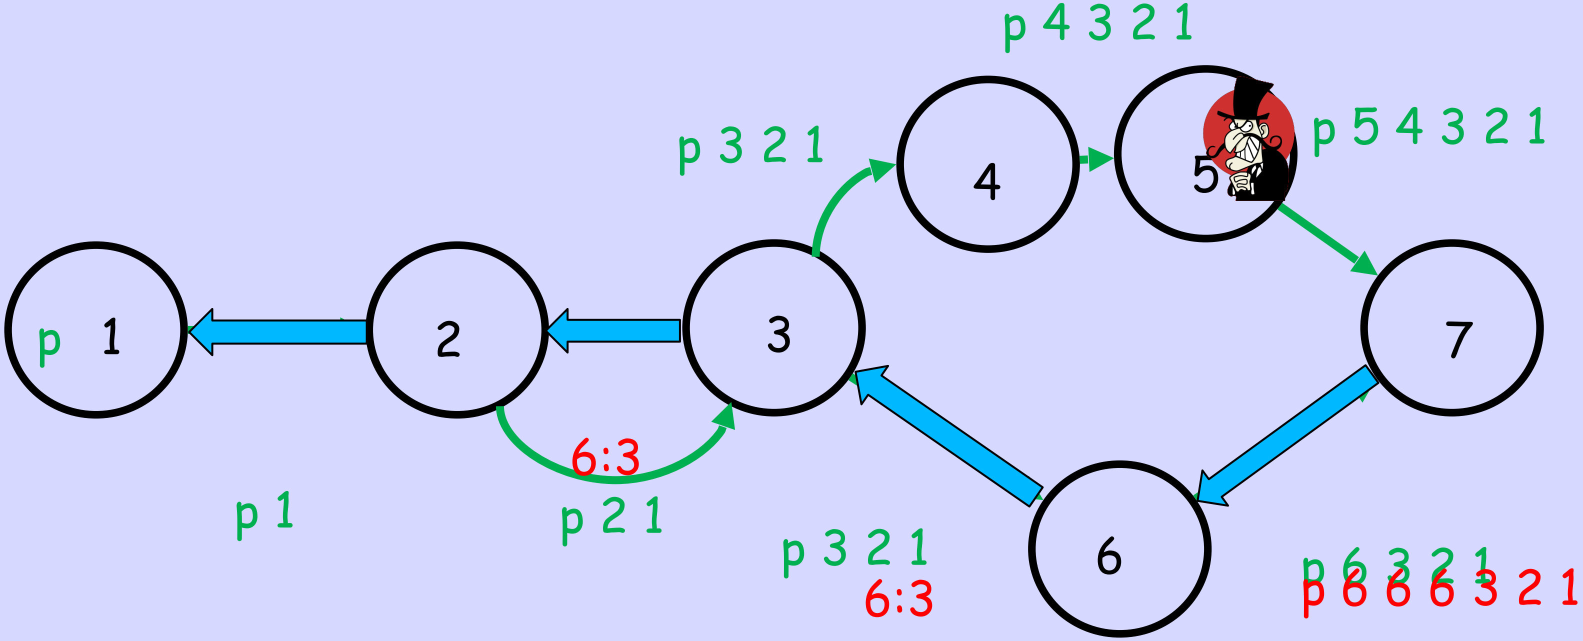

Traffic Steering

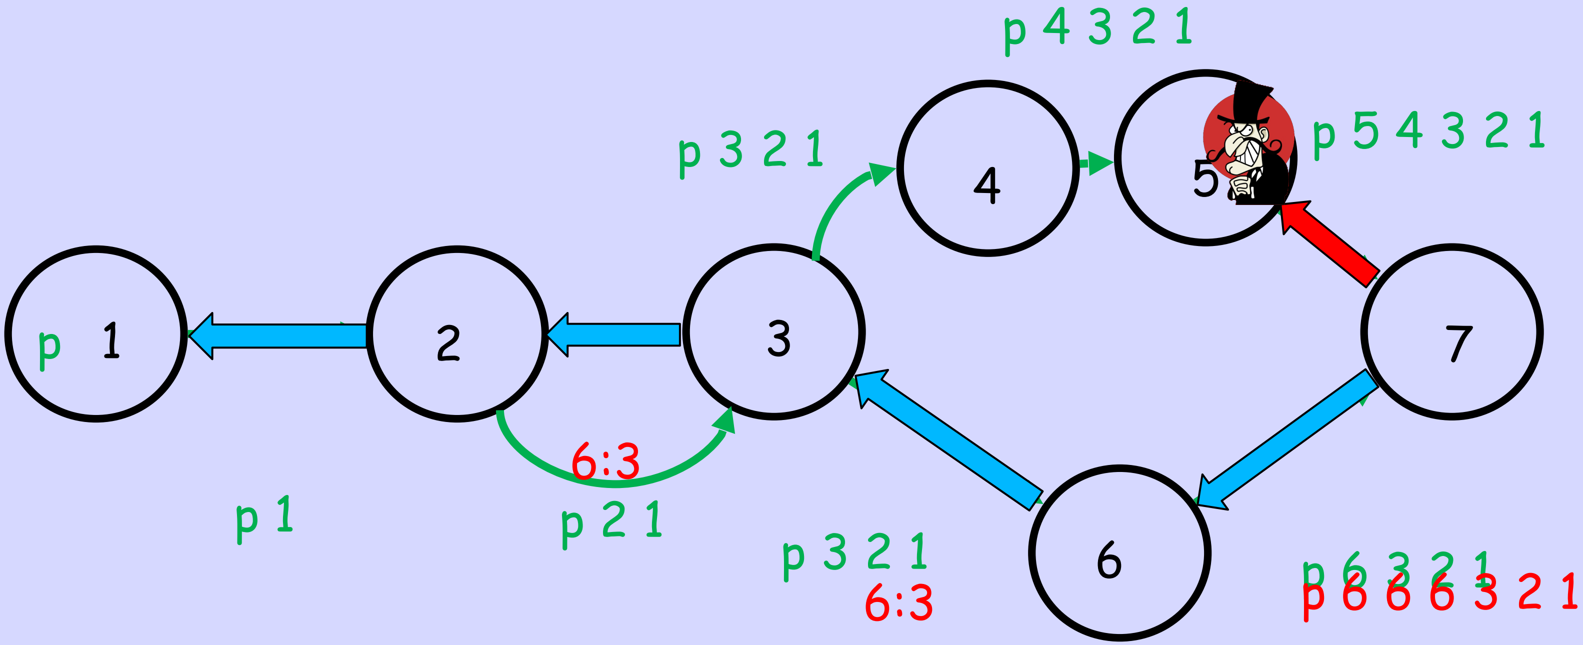

Traffic Steering

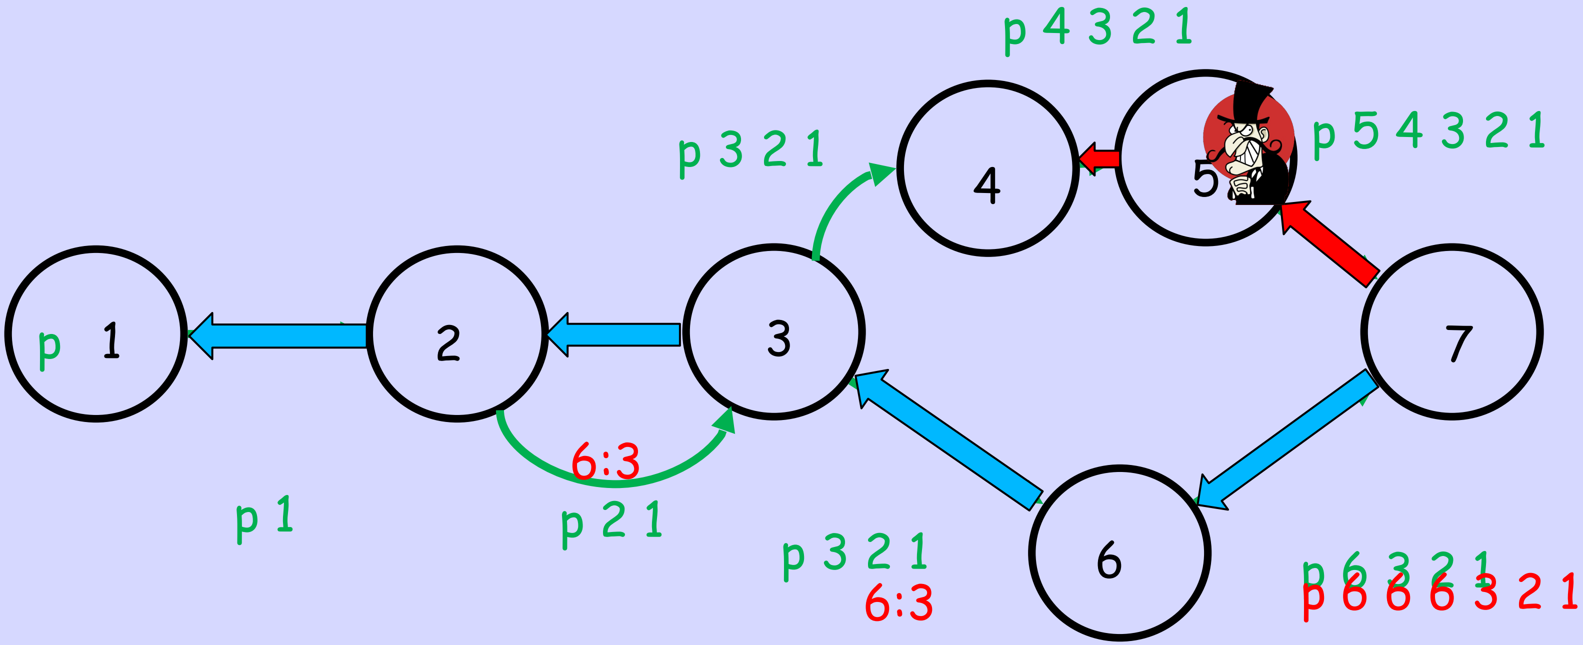

Traffic Steering

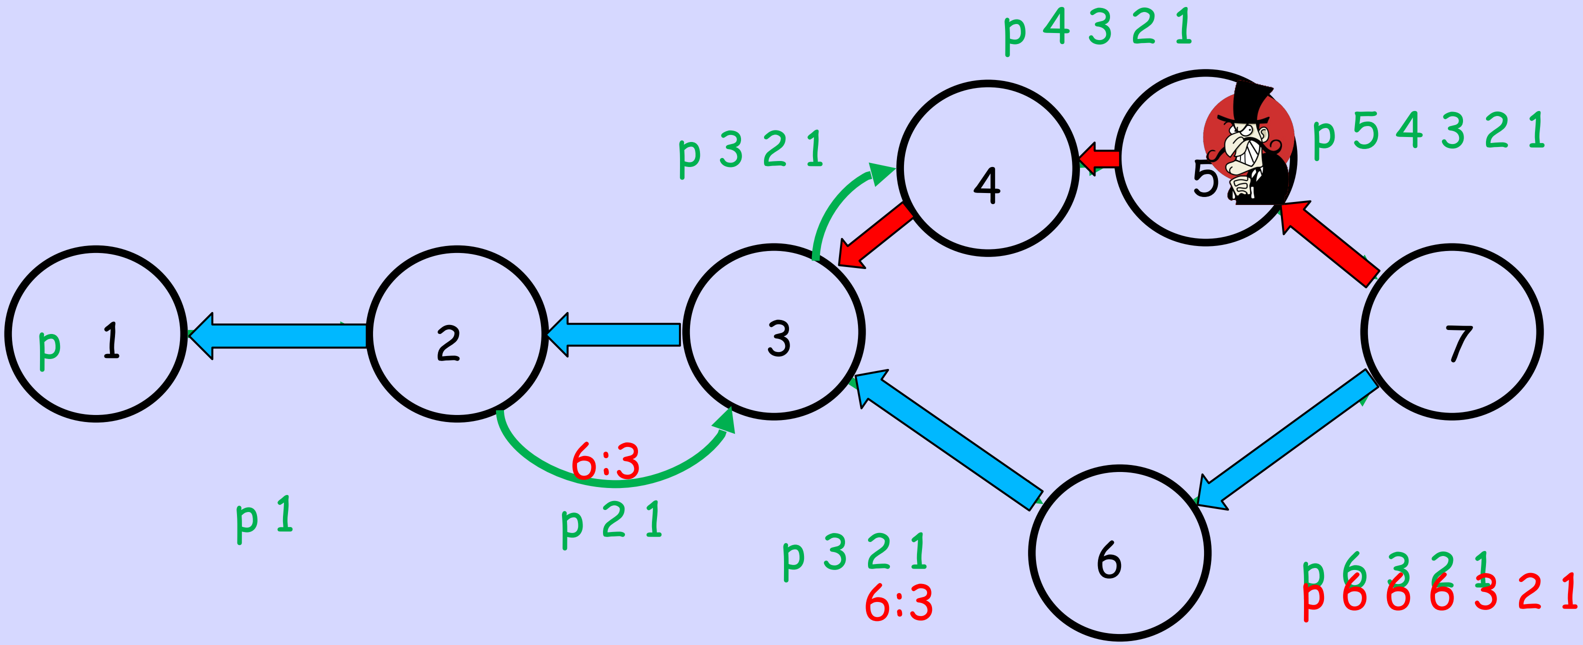

But Is That Realistic?

Yep

https://dyn.com/blog/bgp-dns-hijacks-target-payment-systems/

“BGP hijacks made use of BGP communities to shape route propagation. Although they also changed origins, which was the giveaway.”

It's the Cloud, Man

- ASN value ambiguous: who is "sender", "recipient"
- No defined semantics, values can mean anything
- Used both for signaling and triggering of actions
- No cryptographic protection
- Attribution is impossible
- It is hard to apply filters or understand what is going on

I Read it on the Internet

- Communities can be modified, added, removed by every AS
- No attribution is possible
- No cryptographic protection
- Yet operators bet on their 'correctness'
- Large communities partially improve the situation

Don't Propagate Without Thinking Very Deeply

- On Input – Drop anything not addressed to you, unless special agreement
- On Output – Drop everything except signals from you to the direct peer
- And Beware Cisco 'mis-feature' re well known communities

draft-ietf-grow-wkc-behavior-00

Design on a Serviette

Die by Serviette

ONLY
YOU
CAN PREVENT
WILDFIRES
SMOKEY
Ad Council
SMOKEYBEAR.COM
SMOKEY